

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

From: pieterh@sci.kun.nl (Maintainer)

Newsgroups: comp.sys.ibm.pc.hardware.storage, comp.sys.ibm.pc.hardware.misc

Subject: Enhanced IDE/Fast-ATA/ATA-2 FAQ [1 of 2]

Date: 24 Jan 2000 10:58:49 GMT

Message-ID: <86hb99\$70e\$1@wnnews.sci.kun.nl>

Summary: This FAQ addresses issues surrounding Enhanced IDE, ATA-2, ATAPI and Enhanced BIOSes. It includes practical questions, background information and lists of net resources.

Precedence: bulk

Name: Yet Another Enhanced IDE/Fast-ATA/ATA-2 FAQ

Version: 1.92

Archive-name: pc-hardware-faq/enhanced-IDE/part1

Posting-Frequency: Monthly (the 24th)

Last-modified: 1998/01/23

URL: <http://come.to/eide>

Maintained-by: Peter den Haan <pieterh@sci.kun.nl>

The Enhanced IDE/Fast-ATA FAQ

John Wehman (jwehman@got.net) and Peter den Haan

(pieterh@sci.kun.nl)

v1.92, 1998/01/23

This FAQ addresses issues surrounding Enhanced IDE, ATA-2, ATAPI and Enhanced BIOSes. It includes practical questions, background information and lists of net resources.

1. First off...
2. Introduction
 - 2.1. EIDE and Fast-ATA
 - 2.2. IDE and ATA
 - 2.3. ATA-2
 - 2.4. ATA-3
 - 2.5. ATA Packet Interface
 - 2.6. Ultra-ATA

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

2.7. ATA-4

2.8. The Enhanced BIOS

2.9. The secondary port and beyond

3. General questions

3.1. What are the main features of EIDE and Fast-ATA?

! 3.2. What are the main features of Ultra-ATA?

! 3.3. Are those rumors about buggy interfaces true?

! 3.4. What is a megabyte?

4. Before buying...

4.1. What should I look for in an ATA-2 (Fast-ATA, EIDE) interface?

4.2. Is my BIOS Enhanced? How can I recognize this?

4.3. Is my 50MHz VL bus system compatible with ATA-2 interfaces?

! 4.4. Are old IDE drives compatible with EIDE and U-ATA interfaces?

! 4.5. Can I use EIDE and U-ATA drives with my old (E)IDE interface?

4.6. Will an ATAPI (EIDE) CD-ROM work with an IDE interface?

4.7. Do I need a BIOS update to connect an ATAPI CD-ROM?

4.8. Can large drives be used with my old BIOS?

4.9. I need an Enhanced BIOS to access >504MB. How do I get one?

4.10. Can I use a large harddisk with OSs other than DOS/Windows?

4.11. I heard that Win95 provides support for large IDE disks.

4.12. Can I use more than 2 or 4 hard drives?

4.13. Does <insert operating system> support ATAPI (IDE) CD-ROMs?

4.14. I need a PCI ATA-2 interface that uses only one slot.

4.15. Will adding an Ultra-ATA interface help my drives' performance?

5. Installation

5.1. Which drive/device should be Slave?

5.2. Does it matter how I connect the devices to the cable?

5.3. Does an old HD or CDROM slow down a new drive?

5.4. I need a longer IDE cable; how long can I make it?

5.5. Can I safely move my harddisk between computers?

6. Configuration

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

- 6.1. Should I use my interface's device drivers?
- 6.2. What should I use: Normal, LBA or Large?
- 6.3. Can I safely change the BIOS' relation mode (None, LBA, Large)?
- 6.4. Should I enable EIDE busmastering in my BIOS?
- 6.5. FDISK sees only 504MB of my disk!
- 6.6. FDISK will partition only 2GB.
- 6.7. I have no fancy EBIOS, but I have an 1GB partition and it works.
- 6.8. I have software to break the 504MB barrier. Should I use it?
- 6.9. Can I keep my Disk Manager/EZDrive partition with a new BIOS?
- 6.10. I need to remove Disk Manager/EZDrive.
- 6.11. I can enable write caching on my harddisk. Should I?
- 6.12. My drive letters have changed!

7. Performance

- 7.1. Why isn't my drive as fast as it's advertised to be?
- 7.2. Why is my new EIDE, U-ATA interface no faster than my old one?
- 7.3. I benchmarked my drive and it's getting slower!
- 7.4. A guy on the net has the same drive and he gets 33MB/s!
- 7.5. Busmastering drivers don't do much...

8. Troubleshooting

- 8.1. Why do I get NO ROM BASIC, SYSTEM HALTED?
- 8.2. I have problems with my Award 4.50G BIOS and large drives.
- 8.3. My 2GB+ drive is not recognized or crashes.
- 8.4. My 4GB+ drive has problems in MSDOS 6.22 or below.
- 8.5. Help! I have OnTrack Disk Manager installed and now...
- 8.6. My mode 3-4-5-6 drive has data corruption problems.
- 8.7. Are there supposed to be bad sectors on the drive?
- 8.8. I keep getting CRC errors and serial overruns.
- 8.9. Can my drive do 32-bit access?
- 8.10. Help! Windows 3.x 32-bit disk access doesn't work anymore!
- 8.11. Help! Windows for Workgroups' 32-bit file access fails!
- 8.12. Help! Win95 indicates my drive uses compatibility mode!

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

- 8.13. My partitions become a mess in MS-DOS mode.
- 8.14. Win95 sees my partitions, but DOS doesn't.
- 8.15. Win95 keeps losing my second IDE channel!
- 8.16. The Win95 busmastering drivers won't work.
- 8.17. My ATAPI (IDE) CD-ROM driver won't recognize the CD-ROM.

9. The user's net.resource guide

- 9.1. The user's net.resource guide: I/O card drivers
- 9.2. The user's net.resource guide: hardware manufacturers
- 9.3. The user's net.resource guide: software and BIOS houses

10. ATA: harddisks

- 10.1. How does ATA(-2) work?
- 10.2. What are PIO modes?
- 10.3. What are DMA modes?
- 10.4. How are the ATA(-2,PI) I/O ports assigned?
- 10.5. What does an ATA-2 interface do?
- 10.6. What is Block mode?
- 10.7. What is LBA?
- 10.8. How does security work?
- 10.9. What is S.M.A.R.T.?
- 10.10. What is PRML?
- 10.11. What are MR heads?

11. ATAPI: CD-ROMs and tapes

- 11.1. How does ATAPI differ from, and coexist with, ATA(-2)?
- 11.2. What's so special about the secondary port?

12. The EBIOS: translation

- 12.1. Why translation?
- 12.2. How does translation work?
- 12.3. I'd like to know how translation works in detail.
- 12.4. What is in the Enhanced Disk Parameter Table?

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

12.5. How many types of translating/Enhanced BIOSes are there?

13. Software details

13.1. Details on OnTrack Disk Manager

13.2. How does Windows' 32-bit disk access work?

14. Hacker's resource guides

14.1. The hacker's documentation guide

14.2. The hacker's net.resource guide

1. First off...

This FAQ is the work of John Wehman (jwehman@got.net) and Peter den Haan (pieterh@sci.kun.nl). The homepage of the FAQ is

- o <http://come.to/eide>

The HTML version should also be online on

- o <http://www.wi.leidenuniv.nl/ata/>

and the text version is available by FTP from

- o ftp://ftp.netcom.com/pub/cl/clau/ide_ata/

- o <ftp://ftp.rahul.net/pub/lps/hard-disk/>

- o <ftp://ftp.wi.leidenuniv.nl/pub/faqs/>

- o <ftp://rtfm.mit.edu/pub/usenet/news.answers/faq/enhanced-IDE/>

You can also get it by e-mail from pieterh@sci.kun.nl by sending a

message with "EIDE FAQ text" in the Subject: header. The body of

the message will be ignored. You can replace "text" by "PostScript"

or "html" if you want something more fancy than plain text.

Anything else will probably break the mail server and cause it to

send you an uuencoded coredump :-)

DISCLAIMER.

The information in this FAQ comes without any warranty. The

authors or distributor will not accept responsibility for any

damage incurred directly or indirectly through use of the

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

information contained in this FAQ.

COPYRIGHT.

This document is copyright (c) John Wehman and Peter den Haan. You are free to distribute it by electronic means and make as many copies as you want on electronic or magnetic media, convert it to a different format, and embed control sequences, as long as the text of this document remains unmodified, with copyright notices intact. Splitting up this document is allowed on the condition that the resultant parts are presented as a recognizable whole that makes up the full, unmodified text. You are allowed to make paper copies of this document for personal use; however, distribution of printed copies for profit is prohibited without the authors' prior permission.

THANKS.

To everyone on comp.sys.ibm.pc.hardware.storage for encouragement and feedback, especially those who took the time to communicate their comments in some detail. To Hale Landis without whom the FAQ would have been considerably less accurate on a number of key points; to Stephanie Brady from OnTrack for extensive Disk Manager related information; to Aron Eisenpress for critically reading the FAQ time and again; and many, many others.

TODO.

- o Drop the definition of megabyte as 1048576 bytes, which is nonstandard in this context, and use 1000000 byte MBs instead.
- o Reorganize the FAQ, with a complete overhaul of the text.
- o Write howto type documents walking through specific tasks.
- o Tweak the HTML output more. Perhaps find another translator.

NOTE.

The text, PostScript and HTML versions of the FAQ are all

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

produced from a common linuxdoc-SGML source. This solution is not perfect and all versions do contain some infelicities. Apologies. Suggestions for improvement are most welcome, as are complimentary copies of more professional multi-format authoring tools.

2. Introduction

The aim of this introduction is to make you familiar with the most important buzzwords used in this FAQ and to provide a concise overview of the issues involved. To get the most out of the information in the FAQ proper, start here.

2.1. EIDE and Fast-ATA

! The IDE field has seen a great number of changes since a few years
! ago. These novelties are commonly known under the names Enhanced IDE
! (EIDE) and Ultra-ATA. EIDE has caused a lot of confusion since it is
! merely a marketing program from Western Digital which, in turn, builds
! on a couple of real standards: ATA-2 and ATAPI. Fast-ATA, launched by
! Seagate and endorsed by Quantum in response to WD's marketing is
! similar. It builds on ATA-2 only. Ultra-ATA is a widely supported
! extension of ATA-2 (and ATA-3) adding high speed DMA modes.

2.2. IDE and ATA

IDE (Integrated Drive Electronics--or numerous other interpretations) and ATA (AT Attachment) are one and the same thing: a disk drive implementation designed to integrate the controller onto the drive itself, thereby reducing interface costs, and making firmware implementations easier. This low cost/easy integration created a boom in the disk drive industry, as PC integrators readily ate up the low-cost alternative. Since the late 80's, ATA (as it is properly called) has become the drive of choice for the cost inhibited buyer.

2.3. ATA-2

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

When it became clear that improvements in harddrive technology and the increasing demands made by software would eventually strain the capacity of the ATA interface, the drive industry, in the form of the Small Form Factor (SFF) Committee, created a compatible extension of ATA called ATA-2. This standard not only adds faster PIO modes and DMA modes, but also improves upon Plug'n'Play and compatibility with future revisions of the standard.

While there is also a new way of addressing sectors on the harddisk (LBA), this is merely a simplification. Contrary to common myth LBA proper has nothing to do with breaking the famous 504MB (528 million bytes) barrier. In fact, even in the old ATA/IDE standard the capacity limit is well over 100GB.

2.4. ATA-3

The latest revision of the standard is ATA-3. Key features include: improved reliability, especially of PIO mode 4; a simple password-based security scheme; more sophisticated power management; and Self Monitoring Analysis and Report Technology (S.M.A.R.T.), allowing the drive to warn you about certain types of impending failure.

ATA-3 does not define any faster modes. Even though some manufacturers were once marketing "mode 5" equipment, there will never be a PIO mode beyond ATA-2 PIO mode 4.

2.5. ATA Packet Interface

One of the disadvantages of ATA is that it was designed for harddisks only. That was fine back when a high end PC shipped with just a floppy drive and a 40MB harddisk, but today CD-ROM and tape drives are commonplace devices that should preferably run off a single low-cost interface. The ATA Packet Interface (ATAPI) is a standard designed for devices such as CD-ROMs and tape drives that plug into an ordinary ATA (IDE) port.

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

The principal advantage of ATAPI hardware is that it's cheap and works on every PC with an IDE or 'EIDE' adapter. ATAPI tape drives can enjoy superior performance and reliability compared to the popular QIC117 'floppy' tape devices.

Beware that although ATAPI devices plug into the IDE interface, they differ considerably from an IDE harddisk. Caching controllers and other intelligent interfaces will not work unless they're ATAPI aware. Booting from an ATAPI CD-ROM is only possible with the latest BIOSes.

2.6. Ultra-ATA

Ultra-ATA is somewhat similar to Ultra-SCSI in the sense that it bridges the gap between the current standard (ATA-3) and ATA-4, which isn't quite finished yet. Ultra-ATA adds a new, high performance mode: DMA/33 with 33MB/s bandwidth, twice that of DMA mode 2. It is supported by Intel's Pentium II chipset, the TX Pentium chipset, the latest Intel competitors and a few add-on (PCI) cards.

2.7. ATA-4

Work on the next standard, predictably called ATA-4, has already started. First and foremost it is a much-needed attempt to merge ATA-3 and ATAPI into one. Faster transfer modes are also on the agenda, of course: it will incorporate Ultra-ATA's DMA/33, at least.

Regarding other features, the future for strong command overlap looks rather bleak: the current proposal is a hideous animal and important players like Microsoft do not plan on supporting it in its present form. More limited forms of overlap stand a somewhat better chance of surviving.

2.8. The Enhanced BIOS

A rather different issue was the 504MB (or, equivalently, 528 million bytes) capacity limit that becomes apparent when accessing IDE drives

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

through the BIOS. It is caused by the disk geometry (cylinders, heads, sectors) supported by the combination of an IDE drive and the BIOS' software interface ('int13'). Both IDE/ATA and the BIOS are capable of supporting huge disks, but their combined limitations conspire to restrict the useful capacity of the drive to 504MB. Since only MSDOS still uses the BIOS for harddisk access, this is sometimes erroneously thought to be an MSDOS limitation; other operating systems experience the same restrictions at boot time though.

An Enhanced BIOS works around this problem by representing the drive to the software using a different geometry than the native geometry of the drive itself. This juggling act is called 'translation'. For example, if your drive has 1500 cylinders and 16 heads, a translating BIOS will make software programs think that the drive has 750 cylinders and 32 heads.

You do not need an "EIDE" adapter to use harddisks greater than 504MB. The de facto standard is described in the Microsoft/IBM "INT 13 Extensions" document. Phoenix has presented a superset in their "Phoenix Enhanced BIOS" specification. Phoenix, AMI, Award and MR BIOSes are based on the Microsoft/IBM specification.

2.9. The secondary port and beyond

Last but not least, the usual limit of two devices was far too restrictive if CD-ROMs and tape drives were to be connected to the IDE interface in addition to the harddisk(s). Fortunately the solution was already known in the form of a so-called secondary harddisk interface. The possibility of such an interface has been there for a long time, but support was lacking.

There is nothing special about a secondary channel; it is an ordinary (E)IDE port that uses a different interrupt and I/O addresses to avoid clashes with the ordinary (primary) one. A secondary interface allows you to connect another ribbon cable with two more ATA devices

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

(harddisk, ATAPI CD-ROM or ATAPI tape). Today, many interfaces combine both primary and secondary port on a single board to make a dual-ported interface that handles up to four devices.

To use harddisks on the secondary port with DOS and Windows 3.x, you will need BIOS (either system BIOS or adapter BIOS) or driver support. You can recognize a BIOS with four drive support by the fact that it allows for four sets of drive parameter in the BIOS setup.

There are two further (semi-)standard channels beyond the secondary port: the tertiary and quaternary ones. Some soundcard IDE interfaces can be configured as tertiary or quaternary. See section 10.4 for the I/O and IRQ assignments. Software support for these is still rare.

3. General questions

3.1. What are the main features of EIDE and Fast-ATA?

The fast transfer modes (PIO modes 3 and 4, multiword DMA modes 1 and 2) are the cornerstones of Fast-ATA and EIDE. These are marketing terms contrived by disk drive manufacturers. Enhanced IDE is a Western Digital trademark; Fast-ATA is a term coined by Seagate, and endorsed by Seagate and Quantum.

EIDE consists of:

- o Fast transfer modes: PIO mode 3 or better, multiword DMA mode 1 or better,
- o LBA mode (explained in section 10.7),
- o Four devices on the ATA interface: secondary port,
- o No 504MB limit with DOS: WD Enhanced BIOS,
- o tape backup and CD-ROM devices on the ATA interface: ATAPI.

Fast-ATA and Fast-ATA-2 embrace:

- o PIO mode 3 (and 4 for Fast-ATA-2), multiword DMA mode 1 (and 2 for Fast-ATA-2),
- o Read/Write multiple commands (also known as block mode; see Q10.6),

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

o LBA mode.

The difference between the two schemes is mainly in the scope of EIDE. Hardware can be tagged 'EIDE' even if only part of the EIDE feature set has been implemented, which can lead to some confusion. This FAQ will avoid the term EIDE whenever possible, discussing its component parts (ATA-2, ATAPI, etc) instead. This allows you to see exactly what 'EIDE' features you need in your specific situation.

! 3.2. What are the main features of Ultra-ATA?

! To everything EIDE has on offer, Ultra-ATA adds the following:

! o Even faster transfer modes, most importantly DMA/33 with double the bandwidth of DMA mode 2 and PIO mode 4. There are no new PIO modes.

!

! o Improved reliability using DMA/16 and DMA/33. A checksum is added to the data sent over the ATA interface. That way, data corruption can be detected and the data retransmitted. On an ordinary EIDE interface you wouldn't notice the corruption until it was too late.

! Ultra-ATA was first proposed by Quantum and is widely accepted now.

! 3.3. Are those rumors about buggy interfaces true?

Very true, unfortunately.

This FAQ doesn't really deal with specific interfaces, but two very popular interface chips have been shown to contain bugs too serious to ignore:

o the CMD640x, a dual-channel PCI to EIDE interface used on many mainboards (Intel!) and interface boards, has a number of dangerous bugs you need to be aware of.

o The PC-Tech RZ-1000, used on AT&T, Dell, Gateway and Intel boards, also has two data-corrupting bugs. See also

<<http://www.intel.com/procs/support/rz1000/index.htm>>.

In both cases, the corruption occurs only in specific software environments and is very subtle; you can go on working for months

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

without suspecting anything more than buggy software. The damage can be immense. For all the details, look at Roedy Green's (roedy@bix.com) "PCI EIDE controller flaws" FAQ included with his EIDE test `<ftp://garbo.uwasa.fi/pc/diskutil/>` program which will test your system for the bugs.

BE WARNED that you're playing Russian roulette with your data if you continue working on an affected machine without taking notice of this problem.

3.4. What is a megabyte?

The word "mega" is an ISO prefix designating a factor 1,000,000. A proper megabyte is 1,000,000 bytes. Because computers use binary technology and like working with powers of two, 2^{20} bytes, that is 1048576 bytes, is also usually referred to as a megabyte.

Which of the two types of megabyte you're dealing with depends on the context. For storage devices, a megabyte usually means 1,000,000 bytes. Some software uses "binary" megabytes, though, and will show a smaller capacity than the drive label says! This includes most BIOSes.

This FAQ uses megabytes of 1048576 bytes throughout.

4. Before buying...

4.1. What should I look for in an ATA-2 (Fast-ATA, EIDE) interface?

There are a number of 'EIDE' features which can be desirable in a new interface. Some features, however, may already be present on your system or be unimportant to you.

- ! o The least a modern interface should provide is PIO mode 3 transfers (up to 11.1MB/s) for drives that support it. It should also be able to use slower PIO modes (0, 1 and 2) to ensure compatibility with older drives.
- ! o True direct memory access (DMA) is found mostly on interfaces

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

! integrated into mainboards, Intel 430*X based boards being the most
! common example. This will improve system performance in
! multitasking operating systems. Ultra-ATA will give you DMA/33,
! which actually isn't all that much faster but a lot safer for your
! data.

- o Since ATAPI CD-ROMs have rapidly become very popular in the low-end market, and ATAPI tapestreamers are similarly taking off, a secondary port allowing you to connect a total of four ATA* devices is hardly a luxury. Note that a few modern soundcards provide a secondary or tertiary ATA/IDE interface instead of the traditional proprietary CD-ROM connectors. Beware of conflicts in combination with a dual-ported interface.

If you intend to connect harddisks to the second port and use them with DOS or Windows 3.x, remember that many older BIOSes have no support for the secondary channel and many interfaces do not ship with the required drivers.

- o You will want an on-card Enhanced BIOS too if your mainboard BIOS doesn't support translation or if its support is buggy or outdated. The interface BIOS will override the mainboard BIOS' harddisk routines.

A BIOS ROM will, just like a mainboard with integrated EIDE, often have the added advantage that you don't need separate DOS drivers. The BIOS on an interface will occupy 8 to 16k of UMB space, though, and you will still want drivers for every other operating system you use.

- o Drivers! Without well-designed drivers most interfaces could as well be old-fashioned ISA cards for all the good they'll do. Even if the card has a BIOS, which usually removes the need for a driver

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

under DOS, you will still need drivers for other operating systems,
! including Windows and Win95. Newer versions of Windows (Win95 OEM2
! and beyond) ship with a large amount of drivers.
! Note that some 'EIDE' interfaces which used to be popular were
! slightly to very buggy. This ranged from minor problems with ATA-2
! compliance to obscure things like the use of a single buffer for both
! primary and secondary channel or a badly designed prefetch buffer,
! both of which may cause data corruption under very specific
! circumstances.

4.2. Is my BIOS Enhanced? How can I recognize this?

Good question.

Unfortunately, I know no easy answer. The mere ability to specify more than 1,024 cylinders in the BIOS setup is not conclusive. In your BIOS setup, drive related settings like "LBA", "ECHS" or even something silly like "Large" are telltale signs of a BIOS with translation support, which should be good for disk capacities of up to 8GB. A copyright before 1994, on the other hand, reduces your chances to something close to zero :-)

- o For AMI, I only have reliable information on their HiFlex BIOS; it can be recognized by its characteristically funky orange and green color scheme. There at least two other types: WinBIOS with a Windows like interface, and a custom BIOS used with Intel Pentium boards. Both may translation even if they have a (much) earlier copyright. I have no further information on those.

AMI HiFlex BIOSes dated 7-25-94 and later and support translation. The date is embedded in the long number displayed at the bottom of the screen on bootup; it must be 072594 or later to support LBA.
40-0100-00101111-111192-486-ABC-F (111192 will not support LBA)
50-0100-001292-00101111-072594-ABCDEF-F (072594 will support LBA)

- o Award seems to call all its BIOSes 4.50G :-) Some have translation support, some don't, some have buggy support (see 8.2). With a BIOS

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

dated 12/31/1994 or later you have full translation support, with one dated earlier (7/29/1994, perhaps earlier than that as well) buggy support only.

- o Phoenix BIOS v4.03 and later are reported to support translation. Some revisions of Phoenix v4.03 may not support it, though, depending on the computer or system board manufacturer.
- o MR BIOS incorporated CHS translation--roughly the same as today's Large mode--as early as 1990!

Note that only BIOSes fully implementing the IBM/Microsoft/Phoenix standards will allow access to disks larger than 8GB. Fortunately, these are becoming more and more common these days. We are not aware of utilities that will detect the presence of such a BIOS.

Western Digital has a utility available by ftp that examines the Enhanced Disk Parameter Table (EDPT), if present. If it finds one on your system, your BIOS has all the bells and whistles to go up to 8GB <<ftp://ftp.wdc.com/drivers/hdutil/>>. Beware that BIOSes conforming to the WD Enhanced BIOS specification won't build an EDPT unless LBA is enabled. Beware also that with a WD EBIOS it is highly unsafe to enable LBA on a disk that already contains data. This is due to a flaw in the specification.

4.3. Is my 50MHz VL bus system compatible with ATA-2 interfaces?

It should be, but there are problematic cases.

It has been observed that some controllers will base their I/O cycle times on the bus clock of the computer. This means that with a 50MHz bus, the cycle time will be faster than with a 33MHz bus. This could lead to undesirable results if these cycle times are faster than the drive can handle. Make sure your controller supports multiple bus speeds in the software setup, and if necessary program it correctly.

4.4. Are old IDE drives compatible with EIDE and U-ATA interfaces?

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

This should work fine.

Older drives do not support the high-speed advanced transfer modes defined by ATA-2 and, later, Ultra-ATA. To retain compatibility with these drives, ATA-2 and Ultra-ATA interfaces have programmable timing. The driver or BIOS queries the drive regarding its maximum transfer rate and will configure the interface to match. A few interfaces are jumper configurable in which case the responsibility is yours.

If you intend to put an older drive on the same cable as an ATA-2 drive, you may want to verify that your interface is capable of using independent timing for master and slave device. Otherwise, your old drive might end up slowing down your newer one because the controller is forced to use the lowest common speed. Modern interfaces like the Intel PIIX3 (430HX and VX and 440FX chipset), PIIX4 (430TX chipset) and, to a lesser extent PIIX (430FX chipset) don't suffer from this problem.

4.5. Can I use EIDE and U-ATA drives with my old (E)IDE interface? Fast-ATA and EIDE devices can be used without fear on regular ISA or VLB IDE controllers. The thing about advanced modes is that the drive is ready for the data in the fastest mode; the slower the controller, the easier it is on the drive. Of course you should expect the transfer rates to come out lower than on a full ATA-2 interface.

! Likewise, an Ultra-ATA drive will do fine on an ATA-2 (EIDE) interface. Better even, using DMA mode 2, its performance will be hardly any slower than it would be on an Ultra-ATA interface.
!

It is not necessary to use an EIDE interface for harddisks of over 504MB. An Enhanced BIOS is all you need for DOS and Windows; for many other operating systems such as OS/2 or Linux you won't need even that.

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

4.6. Will an ATAPI (EIDE) CD-ROM work with an IDE interface?

Yes. These devices were designed to be compatible with ATA (IDE) from the very start, and should work fine.

4.7. Do I need a BIOS update to connect an ATAPI CD-ROM?

No. These CD-ROMs ship with a driver that provides complete support. The BIOS doesn't need to support them and usually doesn't even know they're there. Some modern BIOSes are ATAPI aware and have a special setting; barring those, you can generally leave the corresponding harddisk entry at 'Not Installed'.

4.8. Can large drives be used with my old BIOS?

Yes, in principle, but there's still the small matter of the 504MB barrier. If the new drive exceeds this capacity, you must have an Enhanced BIOS in order to use more than 504MB with DOS, Windows or Win95 (see Q4.10 for other operating systems). It is important to note that this is a software only issue: you do not need a special interface to support large drives.

BE WARNED that some older BIOSes will incorrectly handle drives with more than 1024 cylinders. Instead of truncating the number of cylinders to 1024, they use the cylinder count modulo 1024, which means that a 2300 cylinder drive will appear to have just 252 cylinders. The solution is to specify 1024 cylinders in the BIOS setup.

This becomes a real problem when a BIOS with this error also won't allow you to enter the drive parameters manually. Among IBM systems, the 1991-vintage PS/2 models 35sx and 40sx, the ValuePoint I, and at least some of the ValuePoint Si models have this flaw. With these machines, you should either try go get a BIOS upgrade from <ftp://ftp.pcco.ibm.com> if you have a flash BIOS, or verify that the drive you intend to use has a way to fake having no more than 1024

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

cylinders (Western Digital has a utility for their 540MB drive, and many Maxtor drives have a jumper). The Promise EIDEMax is also reported to work, and will allow you to use drives of up to 8GB in full. Note that some operating systems, such as Linux and OS/2, allow you to override the BIOS' drive geometry information, which may also help addressing this problem. Another option is using OnTrack Disk Manager version 7.

4.9. I need an Enhanced BIOS to access >504MB. How do I get one?

The options are: a new mainboard BIOS, add-on BIOS, or software.

- o To upgrade your mainboard BIOS, either with a newer version of your current BIOS or using a custom made BIOS from firms such as Microid Research (MR BIOS), is probably the best option. It may or may not be more effective to replace the entire mainboard and get all the latest bells and whistles for \$100 more.
- o The next option is to purchase a card with an add-on BIOS that effectively replaces the harddisk portion of your mainboard BIOS. An example is LBA Pro from Storage Technologies, sold in the USA by Unicore software, or AMI's AMIDisk Extender. DataTechnology (DTC) also sell such a card.
- o Most BIOSes on ATA-2 interfaces, if they have a BIOS at all, also do this. That may be a nice way to go if your present interface isn't up to snuff. For the ISA bus, there are a few "EIDE" interfaces that have a BIOS: examples are the Promise EIDEMax and the SIIG IDE Enhancer.

Either way, you will have to repartition and reformat the (large) harddisk.

- o A final option is to use a software solution, like Disk Manager or EZDrive. Software is often supplied free with the drive and relatively easy to install, but has inherent drawbacks. See also

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

Q6.8.

4.10. Can I use a large harddisk with OSs other than DOS/Windows?

Using a large harddisk is usually no big problem, even if you don't have an Enhanced BIOS. However, some OSs don't understand translation, which makes the combination with DOS, Windows and Win95 problematic.

With operating systems such as NetWare, Unix, Win/NT and OS/2, the only thing you need to use large ATA disks is a BIOS that allows more than 1024 cylinders in the drive type setup. There is one caveat though: the BIOS is still used to boot the operating system, so you will have to ensure that everything necessary to get the OS running in the first place resides below cylinder 1024. Remember that if you have an Enhanced BIOS, drives up to 8GB will appear to have no more than 1024 cylinders, so in those cases these boot restrictions are removed.

Once running, these operating systems use their own software to control the disks (Win95 also does this, but has a special position; see below). That way, they are not subject to the BIOS' restrictions such as the capacity limit. Unfortunately, this also means that if you have a translating EBIOS or software driver, the OS has to be aware of the translation scheme used, or conflicts will arise between the operating system and DOS/Win/Win95. If you can set up partitions so that all DOS and boot partitions reside below the first 1024 cylinders (504MB), you can avoid translation altogether and all the hassles with it.

Older operating systems don't understand translation at all. Newer ones (OS/2 3.x, Linux 1.2 or better, Win95, SCO 5.0.x) will handle standard translation schemes out of the box, but not always those employed by some software drivers (EZDrive, Disk Manager v6.x or older). In the case of DM 6, fixes or updates are available for some operating systems (see above).

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

Novell has a NetWare driver IDE.DSK version 3.0 dated September 2, 1994. This version of the driver uses the Identify Device ATA command to get the drive parameters and ignores the BIOS parameters. This means that Novell now works with big IDE drives.

A final remark: OS/2 enforces DOS compatibility for FAT partitions. That means that without an Enhanced BIOS, only HPFS partitions can extend beyond cylinder 1024.

4.11. I heard that Win95 provides support for large IDE disks. True, but it doesn't necessarily mean you can actually benefit from that support in all cases.

Despite a large number of significant changes, in its way of handling harddisks Win95 largely resembles Windows for Workgroups. Just like OS/2, Linux and other operating systems, Win95 uses 32-bit protected mode drivers for the harddisks (unless it's using compatibility mode), and is happy to handle harddrives of respectable size: well over 100GB. However, unlike OS/2 and Linux, and like Windows for Workgroups, Win95 boots from the same old DOS we all love to hate. Even once Win95 has booted, DOS hasn't gone. Win95 always keeps it behind the scenes and uses it to run devices in compatibility mode. In 'safe mode', important for troubleshooting, Win95 completely relies on DOS.

What all of this implies is that even though Win95's protected-mode disk driver may support gargantuan drives, it will never change the setup it inherits from DOS in order to retain compatibility with the latter. If you can only get at the first 504MB of your drive in DOS, it will be the same in Win95. You will still need one of the usual > 504MB solutions mentioned elsewhere in this FAQ. For details, see also Microsoft KnowledgeBase article Q126855, "Windows 95 Support for Large IDE Hard Disks".

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

4.12. Can I use more than 2 or 4 hard drives?

For DOS and Windows you probably want BIOS support for all your harddisks. Older BIOSes support drives on the primary channel only, just two; newer ones add support for the secondary channel for a total of four. Few BIOSes support more than that; MR BIOS is an exception.

If the BIOS supports just two drives and you want more, there's a software solution available: 3drives (see the net.resource guide). Some interface cards also ship with the necessary support. Creative Labs has drivers for CD-ROMs on the tertiary and quaternary port.

Other operating systems such as Win95 and OS/2 support both primary and secondary interface without aid from the BIOS. Linux will support a nearly arbitrary number of interfaces; you'll have to pass kernel parameters to specify the ones beyond the first two.

4.13. Does <insert operating system> support ATAPI (IDE) CD-ROMs?

The current versions of most operating systems support ATAPI CD-ROMs. That includes OS/2 3.0, Linux 1.2 and above, and Win95.

Unfortunately, hardware manufacturers have started designing these CD-ROM units while the ATAPI standard was still evolving, so there are now a couple of versions implemented in the real world. Moreover, the complexity and novelty of ATAPI means that there are some variations even among implementations of the same revision. Writing an OS ATAPI driver that works with all CD-ROMs, therefore, is a daunting task and not all have succeeded equally well. So although these OSs all support ATAPI, they do not work with all ATAPI equipment.

See Microsoft KnowledgeBase Q131499, "CD-ROM Drives Requiring Real-Mode Drivers", for more details on Win95 support. For OS/2, take a look at <http://www.austin.ibm.com/pspinfo/os2hw.html>; for updated

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

drivers use the latest install disk upgrade or
<ftp://ftp.pc.ibm.com/pub/os2_drivers/atapi.zip>.

4.14. I need a PCI ATA-2 interface that uses only one slot.
There is a problem with PCI and multi-I/O or dual-ported IDE interfaces. Such interfaces need two or more predetermined interrupts which cannot be shared with other PCI devices, which is probably impossible with your PCI slots. The usual workaround consists of either integration of I/O functions on the mainboard, or the use of a tiny 'paddle' board that plugs into an ISA slot.

This is not so much a weakness in the design of PCI, but a conflict between PCI's plug'n'play philosophy and the requirement that these interfaces be fully compatible with oldfashioned I/O cards.

Some vendors use PCI interfaces that rely on a proprietary extension of the PCI bus. This is obviously not portable; you often can't even move the card to a different slot in the same machine. Moreover, such extensions may cause compatibility problems since they use PCI signals reserved for other purposes.

4.15. Will adding an Ultra-ATA interface help my drives' performance?
A new interface may or may not help; it is possible to make a rough prediction if a better interface would really speed things up.

Hardware vendors and marketing people would love to see everyone rush
! out and buy the latest generation of 'Ultra-ATA' adapters. To achieve
! this noble goal they tend to juggle with too-good-to-be-true
! performance figures. The relation between this advertising hype and
! the real world is shaky at best.

The main point to remember is: a slow drive is a slow drive no matter
how good the interface is. If the speed at which the drive physically

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

transfers the data to/from the media is the limiting factor in performance, and it often is, the only way to make things go significantly faster is to purchase a better drive. Note that the ! transfer modes supported by modern drives (those 33MB/s figures) have ! little to do with their real-world performance.

! In addition, an 'ordinary' ATA-2 (EIDE) interface already offers ! respectable bandwidth---the fastest ATA-2 transfer mode is ! theoretically 16.6MB/s, which is more than any Ultra-ATA drive on the ! market today can sustain. In addition it usually supports the CPU- ! cycle-saving DMA modes. An IDE interface, on the other hand, makes a ! much better candidate for replacement since it rarely has a bandwidth ! over 2.5MB/s, which is cramped by today's standards, and doesn't ! support DMA.

How to determine if the drive is the bottleneck? You can get a rough idea using Coretest version 3 <<ftp://ftp.rahul.net/pub/lps/disk/core303.exe>>. This version of Coretest gives two performance figures of interest here: the (B)uffered transfer rate which is an indication of the bandwidth between drive and interface, and the (S)ustained rate which is related to the speed of the drive media. If your drive has a small buffer cache, you may have to use the /B16 option to get the correct buffered transfer rate.

Usually you'll find that the first figure exceeds the second by a considerable margin (say, a factor two or more). This means that the physical properties of the drive itself are the bottleneck, and improving the interface speed any further won't help much. The only thing that may improve performance somewhat is using block mode (using either a BIOS option or a driver). Only if the drive throughput starts to approach the interface bandwidth will you have a fair chance that a new interface will have a large impact on performance, provided the drive supports faster modes than the one currently used.

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

This is not an endorsement of Coretest as a disk benchmark; there are more reliable ones around, such as QBench

<<ftp://ftp.rahul.net/pub/lps/hard-disk/>>.

Be sure to read question 7.2 to get a more complete picture.

5. Installation

5.1. Which drive/device should be Slave?

- o A single device on a cable, either a harddisk (ATA) or CD-ROM (ATAPI), should according to the specs never be configured as slave. However, as any ATAPI CD-ROMs come preconfigured as slave, most software works with a masterless slave CD-ROM simply because it saves tech support calls. Moreover some BIOSes have trouble with an ATAPI device jumpered as master. Formally it isn't a valid configuration though; remember this if you have trouble getting the CD-ROM recognized.
- o Two harddrives on one channel should be configured as master and slave, respectively. It usually doesn't matter which is which, but older harddisks may not work in all configurations if they predate the standardization of the master/slave protocol. In that case a degree of experimentation will be necessary. Some Conner drives have an ATA/ISA jumper: ATA is the now-standard protocol, ISA is used by older Conners.
- o Two ATAPI devices such as CD-ROMs and tapes should work fine as long as one device is master and the other slave, just like harddrives.
- o If you need to put a harddisk and an ATAPI device on the same cable, the spec says the harddisk must be the master. Although the reverse will usually work as well in practice, still you'd better avoid the configuration.

5.2. Does it matter how I connect the devices to the cable?

Not much. If you have only one drive on the cable, it is best to put

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

it at the very end, especially when you're using any of the faster modes. For two devices, it doesn't matter where you put the master and the slave, or which end of the cable you plug into the controller. Just take care that you plug them in the right way: the red wire is supposed to correspond to pin 1.

When Plug'n'Play ATA arrives the Cable Select (CS) setting will be used, and all of this will change. Some name brand machines may already employ a similar setup.

5.3. Does an old HD or CDROM slow down a new drive?

This is not necessarily the case. Still, it is generally preferable to connect older drives and CD-ROMs to the secondary channel.

If this is not feasible, or if you're wondering if you should upgrade, a few points.

- o The speed loss usually referred to is in the interface timing, i.e. the speed at which the devices communicate with the computer. This does not necessarily translate into a real world performance penalty.
- ! o This is mostly an issue with older ATA-2 (EIDE) interfaces and some VL IDE ones. If you have an ordinary ISA IDE interface, it can't get any slower.
- ! o All modern interfaces support distinct timing for master and slave. With these, the slow device does not directly affect the fast one.
- ! o Many CD-ROMs support at least PIO mode 3. This is enough to operate most harddisks on the market today near their maximum speed.

You can use Coretest <ftp://ftp.rahul.net/pub/lps/disk/core303.exe> to determine if and how performance is affected; see Q4.15 for a recipe.

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

OS/2 and Unix users have another reason to put slow ATA devices such as tapes and CD-ROMs on a channel of their own. As long as one unit on a given channel is executing a command, the other is inaccessible. A CD-ROM can easily occupy the channel for 300ms that way.

5.4. I need a longer IDE cable; how long can I make it?

Less than 18 inches. In some cases, the limit is no more than 7 inches.

The cable is a pretty weak link in the whole ATA-2 interface. For the fast transfer speeds used in 'EIDE' systems, there aren't enough ground signals; the cable is unterminated and unshielded. Noise is a real problem. All of this applies to ATA-2 (EIDE) systems more than Ultra-ATA systems, since the latter use checksums to ensure data integrity.

!

! For those reasons, you should take the 18" limit specified by the ATA(-2) standard pretty seriously if you want to avoid data corruption and system hangs. Even worse, some dual-channel 'EIDE' interfaces such as CMD640x based ones have a peculiar design employing only a single buffer for both cables. Most signals on the two cables are directly connected with each other: this means that electrically, the cable lengths add up. Take into account that the copper traces on the circuit board are often a couple of inches long as well and you're facing a maximum cable length of roughly 7" per cable if you want to remain within spec.

If you have difficulties fitting everything in with a standard length cable, consider adding a new plug to it or rolling a complete cable yourself. It's not hard to do, or to track down a friend willing to do it, and you can buy the parts in all electronic parts stores. Do use quality parts, work carefully, and watch that length.

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

Disregard the above at your own peril. Noise induced problems usually pop up sporadically, can be very hard to detect and even harder to track down. Not least because they may appear completely unrelated and involve devices other than the harddisk.

5.5. Can I safely move my harddisk between computers?

Transferring a partitioned and formatted harddisk between computers is potentially dangerous if they use different translation schemes (see Q 6.3). This includes moving a disk from an old computer that doesn't support translation to a new one that does. This is really a little known issue. Usually everything goes smoothly, but you would not be the first to be caught by surprise.

WARNING. Do not assume it won't happen to you just because it works out fine most of the time---it has destroyed data, and will do so again.

6. Configuration

6.1. Should I use my interface's device drivers?

Yes. Often these drivers are essential to get any kind of performance out of your interface.

The PIO or DMA mode used when transferring data is determined by the interface card. Some cards have jumpers that determine the speed in hardware; these work in the fast mode from the microsecond you switch on the computer.

Most interfaces, however, are software configurable. At bootup, they default to the slowest possible speed. Somewhere during the boot process, a piece of software belonging to your adapter figures out what kind of transfer rates the drives support and configures the controller chip to match. There are a couple of cases to distinguish:

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

- o Onboard I/O with full BIOS support. The controller is fully configured when your computer boots. You can usually set the desired mode for each harddisk in the CMOS setup. Many modern boards fall in this category.
- o Onboard I/O with incomplete BIOS support. For some unfathomable reason, some mainboards do not support or only imperfectly set up their integrated I/O ports. In that case, you'll have to use DOS or other drivers to get full functionality.
- o Interface card with BIOS. This is similar to the two categories above. The main difference is that these cards don't necessarily have setup screens; in that case, they must use other means to determine the transfer mode to be used. For example, the Promise 2300+ uses a combination of jumpers and a table in ROM containing the parameters for a number of different drives. It may or may not be necessary to use drivers for best performance.
- o Interface card without BIOS. Since there is no way the mainboard BIOS can know how to set up all those different interface cards out there, you must use the supplied device drivers to profit from the fast modes. That is, unless your card is hardware configurable using jumpers, which is quite rare.

Usually, there are drivers for other operating systems as well, such as Windows, Win95, OS/2 and so forth. These serve a couple of purposes.

- o The driver may be necessary to configure the adapter as described above. This doesn't apply to Windows, where the DOS device driver usually has already done that job.
- o The standard drivers built into operating systems don't support all of the advanced features of your interface and drives. Examples are 32-bit transfers, block mode and DMA.
- o Windows only: the standard driver (*wdctrl) that ships with Windows and Windows for Workgroups has some serious restrictions. See Q8.10 for details.

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

In view of this it is rather unfortunate that so often, the drivers supplied with an interface are of mediocre quality.

6.2. What should I use: Normal, LBA or Large?

The difference between the three is this.

- o 'Normal' causes the BIOS to behave like an old fashioned one without translation. Use this if your drive doesn't need it (ie. has fewer than 1024 cylinders) or if you want to use the drive with an operating system that doesn't understand about translation.
- o 'Large' or ECHS or XCHS tells the BIOS to use CHS translation. It uses a different geometry (Cylinders/Heads/Sectors) when accessing the drive than when talking to the software through int13. This type of translation works with all drives.

Note. Some BIOSes have a braindead Large implementation which works only for disks of up to 1GB. Fortunately, all larger disks support LBA.

- o 'LBA' differs from 'Large' in that it uses LBA addressing to access the harddisk. The advantage is that it theoretically is a little faster. The disadvantages are that some older drives don't support it, and it often turns out to be slower, depending on the drive.

WARNING. Some BIOSes change the (translated) geometry if you change from Normal or Large to LBA. The same thing may happen if you transfer a disk that has been formatted on an old, non-LBA computer to a new one that uses LBA. This has destroyed data. Don't let it happen to you.

Section 10 tells more about the differences between these three.

6.3. Can I safely change the BIOS' xlation mode (None, LBA, Large)?

Unfortunately, no. Proceed with care.

While with many BIOSes, the sectors on the disk are addressed in the same order independent of the translation mode, a few use a different

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

type of translation algorithm. The latter type of BIOS will shuffle your data as if it were a deck of cards if you alter the translation mode.

Moreover, BIOSes that conform to the WD Guide may use completely dissimilar drive geometries in the software (int13) interface depending on the translation mode. If this happens it will wreak havoc with your data. This represents a major flaw in the WD EBIOS specification.

In both of these cases, after changing the translation mode, you must repartition and reformat your disk.

6.4. Should I enable EIDE busmastering in my BIOS?

The only reason why you would want to enable this option is that DMA modes are less likely to corrupt data than PIO modes. There will be no difference in CPU usage. (when DMA/33 arrives, the improved bandwidth will be another reason).

Unfortunately, at least one user has reported a drastic decline in drive throughput with DMA enabled. The reasons are unclear, so YMMV.

6.5. FDISK sees only 504MB of my disk!

First and foremost, do you have an Enhanced BIOS? See section 2.8 and Q1 for more details. If you do have an EBIOS, make sure you have enabled translation: usually, either 'Large' or 'LBA'. If you see no such options in your BIOS setup, remember that some types of BIOS offer them only when you tell it to autodetect the drives.

Last but not least, remove all old partitions before trying to create new ones after changing the translation mode.

6.6. FDISK will partition only 2GB.

There's nothing wrong; this is a limitation of the DOS FAT and Win95 VFAT filesystems. You will have to create multiple partitions in order to use the full drive size.

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

This limitation has been addressed in Microsoft's new FAT32 filesystem, currently only available in the Win95 OEM 2 release. It allows giant multi-gigabyte partitions. At the time of writing this release can be sold with new hardware only and is unavailable to ordinary mortals.

6.7. I have no fancy EBIOS, but I have an 1GB partition and it works. Some try to work around the 504MB / 1024 cylinders issue by making a large partition using a friend's computer, Linux' fdisk, or something else. They use it for a day or two, conclude that it works, then post a triumphant article claiming that they found the Solution To Everyone's Problems[TM].

It will work... for precisely 1024 cylinders. The very moment the OS or anything else attempts to write something to cylinder 1025 through int13 calls, the write wraps around to cylinder 0. This cylinder happens to hold some of the most important data structures on the disk: the Master Boot Record, partition table, both FAT copies and the root directory of the first partition. Overwrite these and probably only a specialized data recovery company will be able to salvage your data.

Try it if you must. If you know exactly what you're doing, you can make it work using Win95. Sort of. The first error will be fatal. But please don't post any stories about it, recommending the procedure to everyone. The spectres of their valuable data will come back to haunt you.

6.8. I have software to break the 504MB barrier. Should I use it? This is a convenient option, but there are caveats. The software, sometimes bearing an exotic name depending on the licensee, is usually a version of either MicroHouse's EZ-Drive or

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

OnTrack's Disk Manager. Disk Manager, when used on the boot drive, has to resort to some trickery in order to be loaded very early during the boot process (which is necessary for technical reasons). This is accomplished by modifying the Master Boot Record (MBR), the first piece of code the BIOS loads and executes when the computer boots, and storing a Dynamic Drive Overlay (DDO) on the very first disk track. EZ-Drive works in a similar fashion.

An annoying side effect of using a software solution is that operating system installations, which often overwrite the MBR, will render the contents of your harddisk inaccessible. You will need to restore the MBR from the installation floppy to regain access to your partitions.

Moreover, such software tends to create partitions quite different from 'standard' translation schemes as used by most Enhanced BIOSes. Many device drivers dealing with the disk will fail even if they work fine with other schemes. Important examples are (E)IDE interface drivers; remember that without these drivers an interface will in general be much slower. You'll need drivers that are specifically aware of the translation software you use. Also, many operating systems other than DOS will not be able to access or use the drive, at least not 'out of the box'. Disk fixing utilities may fail to work if the partition table or the overlay is damaged.

This also makes it difficult to upgrade to a BIOS based solution. OnTrack Disk Manager version 7 includes a migration utility for the purpose and improves on version 6 in a number of other ways. It allows multiple operating systems and is compatible with most interfaces and drivers, including those for ATAPI CD-ROMs. Owners of version 6.03, which is still often included with hard drives, can download an update patch from OnTrack <<http://www.ontrack.com/pub/software/dmpatch.zip>> that gives some of the benefits of version 7. View <<http://www.ontrack.com/dm.html>> for general information.

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

For older versions of Disk Manager, IBM and Microsoft have fixes for OS/2 (in FixPak 5 or later, or out of the box in Warp FullPack and Warp Connect) and NT (Service Pack 2). Win95 should support Disk Manager and EZDrive out of the box (see Microsoft KnowledgeBase article Q126855, "Windows 95 Support for Large IDE Hard Disks"). More about Disk Manager in section 13.1.

6.9. Can I keep my Disk Manager/EZDrive partition with a new BIOS?

If you have a Disk Manager or EZDrive partition and are upgrading to a translating BIOS, you have three options.

- o The utility is compatible with LBA translation, or has a migration feature. This makes for a smooth upgrade path. (I'm only certain about Disk Manager v7).
- o The above isn't true and you don't want to go through a backup, repartition, restore cycle. Tell the BIOS to use Normal mode for the harddisk and the utility should continue to work as usual.
- o Backup, repartition, restore. See 6.10 if you experience difficulties removing the utility.

Always back up your data before doing a system upgrade.

6.10. I need to remove Disk Manager/EZDrive.

This software is usually installed in the boot drive's Master Boot Record (MBR). Normal repartitioning and reformatting of the drive usually does not refresh the MBR, which can make for a frustrating experience. Fortunately, it's not very hard once you know how.

The best way to deinstall is to follow the procedure outlined in the utility's documentation. If this is not available, the following procedure usually works: boot from a clean floppy with at least DOS, FORMAT.COM and FDISK.COM on it. Then type FDISK /MBR. This should refresh the code in your MBR. After that, repartition and reformat as

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

usual.

For reasons I do not understand, some Disk Manager versions are reported to cling to life rather tenaciously (perhaps due to overambitious virus protection by the BIOS?). In that case, you need DM.EXE. Type DM /Y-. If that fails as well, the following procedure was reported to work by Mark Brown (mrkbrown@netcom.com).

1. run DM (in this case, v.6.03)
2. press ALT-T
3. select (D)isk Sub-System Overview
4. select appropriate hard drive
5. press CTRL-F10 to clear out the MBR
6. press Y to confirm
7. press ESC to exit out of DM, rebooting from a clean floppy

If you have faced similar situations and can add to this, please share your knowledge. A good candidate would be a utility that simply zaps the partition table---any takers?

WARNING. Disk Manager and EZDrive partitions differ from those created by a translating BIOS. Expect your data to be inaccessible after this operation. An exception is v7.x of OnTrack's Disk Manager; DM.EXE has a Migrate feature that works with many BIOSes.

WARNING. Some controllers and security software stores information on track 0 which FDISK /MBR will clear. In this case, the data on the disk is lost anyway, but there are a number of circumstances where this command can destroy data.

6.11. I can enable write caching on my harddisk. Should I?

Do not confuse this with Smartdrive (or whatever) lazy writes: what is meant here is altering the drive's buffer cache management algorithm. This is possible using newer versions of Drive Rocket, with hdparm

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

under Linux, and probably other utilities too. Provided, of course, the drive supports this feature. Sometimes it can also be done using jumpers on the drive.

There seem to be problems with this, if a program will issue a soft-reset (which on Intel Pentium Motherboards also issues a hardware-reset) as soon as it sees the last IRQ, which overall ends up corrupting data. Use with care, and backup.

6.12. My drive letters have changed!

MS-DOS assigns drive letters as follows.

- o Letters a: and b: are reserved for floppy drives.
- o All primary partitions on all (system and adapter) BIOS supported harddisks get their drive letters starting from c:, in order. Normally, you can have just one primary DOS/Windows partition on every drive.
- o Only then, all logical drives inside extended partitions get their letters. This means, for example, that if you had one drive with c: and d:, adding a second drive with one primary partition on it will bump the former d: partition up to e:. If you want to avoid this, do not define primary partitions on all drives except the first one.
- o After that, MS-DOS parses the CONFIG.SYS and AUTOEXEC.BAT files. Some devices such as CD-ROMs have no BIOS support and get their drive letters only here.

Remember that some programs, including disk compression software, may do shuffling tricks with drive letters. Be also warned that some BIOS setup screens, in a misguided attempt at user friendliness, refer to harddisk units using drive letters. In trivial cases, this may be right, but in nontrivial setups the BIOS may be all wrong.

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

7. Performance

7.1. Why isn't my drive as fast as it's advertised to be?

Good question. The basic answer is that the advertised modes are theoretical transfer rates.

This is the case at two levels. First and foremost, the oft-quoted rates do not represent the speed at which the drive can actually read data from, or write data to, the magnetic media. Instead, they give the speed at which data can be exchanged between the drive's buffer cache and the CPU. While the latter gives the more imposing figures, the former has greater impact on real world performance. "It is really as if the government had had a speed limit of 250 km/h on the highways, then raised it to 1600 km/h and tried to impress you by telling you that now you can drive faster"--Aaron Bilger (bilgerar@mentor.cc.purdue.edu).

Second, even once you accept that these transfer rates can be achieved only when the drive happens to have the data ready in the buffer cache, these figures are pretty optimistic. Realistically, drives do more than just give data to the host out of the cache. For each sector transferred to the host, the drive's controller needs to get one from the media; internal controller processing, table updates, positioning and buffer cache management all take some of the controller's attention. All reduce the throughput from the cache to the host.

On top of that, depending on the benchmark used to determine the 'throughput', the rate can vary from 3MB/s to 30MB/s and upwards, all on the same drive. This depends on what the utility actually measures, how it measures it, and even where on the drive it measures it (different zones on the same drive can vary up to a factor two in speed). Plus, system configuration (MHz CPU, RAM, harddisk cache, processor cache) make a difference as well.

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

Bottom line is, whatever benchmark you use, you will not 'see' the advertised transfer rate. The real test is how well it improves your day to day applications. The rest is just fluff.

7.2. Why is my new EIDE, U-ATA interface no faster than my old one?

! There are several reasons why it is possible that an Ultra-ATA
! interface is no faster than an older one even if the drivers are
installed and everything works properly. The first, and most
important, is that the bandwidth of the fastest ATA-2 mode, 16.6MB/s,
exceeds the real world transfer rate of most drives on the market
today by a considerable margin. If this is the case for your drive,
increasing the interface bandwidth beyond this will have no noticeable
effect.

! A further possibility is that your drive does not support the Ultra-
! ATA transfer modes at all. In that case, an Ultra-ATA interface may
! improve performance relative to an old-fashioned ISA bus or VL bus
! interface, but it will be no faster than a good PCI ATA-2 (EIDE)
! interface.

7.3. I benchmarked my drive and it's getting slower!

Believe it or not, but this is completely normal. First, filesystem
fragmentation affects some benchmarks; try defragging the drive.
Second, not all parts of the drive are equally fast.

Physically, a harddrive consists of one or more rotating platters,
where the tracks are concentric circles on these platters.

Obviously, the outermost tracks are longer than the innermost ones.
Because they are longer, they can hold more sectors. As you work your
way inwards and the track length decreases, the number of sectors
decreases in a number of steps. This is referred to as Zone Bit
Recording (ZBR).

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

Back to the benchmarks. Since the platter spins at a constant rate, more sectors in a track give a proportionally higher transfer rate. The very first cylinder of your drive is right at the edge of the platter, in the fastest zone. This is the area that was tested when you got your drive and tried to find out how well it performed. As your drive fills up, you start using higher cylinder numbers---and slower zones. Depending on the type of benchmark you use, this may be reflected in lower scores.

The difference in sectors per track (and hence transfer rate) between the fastest and the slowest zone may be as much as a factor two. Typical drives have anything from five to twenty zones, all with a different number of sectors per track.

7.4. A guy on the net has the same drive and he gets 33MB/s!

"Usenet is like a herd of performing elephants with diarrhea--massive, difficult to redirect, awe-inspiring, entertaining, and a source of mind-boggling amounts of excrement when you least expect it."

-- Gene Spafford (spaf@cs.purdue.edu), 1992

7.5. Busmastering drivers don't do much...

First of all, busmastering will generally not do much with the transfer rates themselves. What it does is relieve the CPU from the chore of shovelling data around, so that it can do something more useful. However if you tend to do only one thing at a time, or you're running a benchmark, the extra CPU time can't be used for anything and busmastering will have no appreciable effect.

Second, you will not see any decrease in CPU usage with the System Monitor in Win95 because it tracks the amount of time spent in the "system idle" loop. However, with a pending I/O operation, the system doesn't spend its spare cycles there, but in the "waiting for I/O to

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

complete" loop.

! Third, you need to use DMA to profit from the improved data integrity

! and enhanced transfer rates of Ultra-ATA.

8. Troubleshooting

8.1. Why do I get NO ROM BASIC, SYSTEM HALTED?

This should get a prize for the PC compatible's most obscure error message. It usually means you haven't made the primary partition bootable or, in Microsoft-speak, 'Active'. Use FDISK to fix this. Don't fret, you won't have to repartition or reformat anything unless you have no primary partition at all.

The earliest true-blue PCs had a BASIC interpreter built in, just like many other home computers those days. Even today, the Master Boot Record (MBR) code on your harddisk jumps to the BASIC ROM if it doesn't find any active partitions. Needless to say, there's no such thing as a BASIC ROM in today's compatibles, and this action ends in the above error message.

8.2. I have problems with my Award 4.50G BIOS and large drives.

If your 4.50G BIOS is dated earlier than 12/13/1994, the address translation table is faulty. To access drives with more than 1024 cylinders, you cannot use address translation (Large) but must use LBA. In versions dated 12/13/1994 or later this has been fixed. Be warned that switching to LBA will probably force you to repartition and reformat your drive; do back up your data.

8.3. My 2GB+ drive is not recognized or crashes.

Many BIOSes dated before 1996 contain bugs making them incompatible with drives of more than 4096 cylinders, which works out to be about 2GB in size. Some BIOSes have similar problems at the 8192 cylinder boundary (4GB). The effects may range from not being able to use the full capacity to a crashing BIOS during bootup or upon detecting the

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

drive.

Possible solutions include

- o a flash BIOS update, if you've got a flash BIOS and your board manufacturer makes a fixed BIOS image available;
- o non-flash BIOS ROM upgrade;
- o a software disk manager like using OnTrack Disk Manager or EZDrive;
- o if the BIOS allows a user definable drive type, manually setting the cylinder count to 4095 (8191). This means you can use no more than the first 2GB (4GB) of the drive.

In the last case, you can use the remaining capacity of the disk to give Linux a test drive.

8.4. My 4GB+ drive has problems in MSDOS 6.22 or below.

Some BIOSes assign a drive of over 8192 cylinders a translated geometry with 256 heads. MSDOS 6.22 and below fail when they try to access the last head.

If your BIOS allows a user definable drive type, use a geometry with 15 heads and 16/15 times the original number of cylinders, rounded down. Thanks to universal translation you can always do this. Remember to write down the geometry somewhere so that you can reproduce it if necessary! If no user definable drive type is possible, there's little you can do about this except upgrade to Win95.

!

8.5. Help! I have OnTrack Disk Manager installed and now...

OnTrack has a web site: <http://www.ontrack.com/>. Most importantly, you can find their Disk Manager FAQ at <http://www.ontrack.com/ontrack/faqhome.html>. There's also a bit about DM in section 13.1. If all of that doesn't help, contact OnTrack tech support at tech@ontrack.com.

Beware that some licensed versions are supported by the OEM rather than by OnTrack. In such cases the OEM usually offers help, FAQs etc.

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

on their web site.

8.6. My mode 3-4-5-6 drive has data corruption problems.

Try using a slower mode or disable fast modes altogether. Mode 3 and especially mode 4 are very sensitive to timing problems, and not all adapters follow the ATA-2 specification really closely. Don't dismiss the possibility too easily: if you changed anything on your system, it is very well possible that a drive which marginally worked so far now starts to corrupt data.

Some controllers seem to configure themselves according to the capabilities of the master drive. This can mean trouble if the slave handles only slower modes.

Moreover, check your cables, and ensure they aren't too long (see Q5.4). Removable drive brackets may also cause problems with fast PIO modes for roughly the same reasons.

8.7. Are there supposed to be bad sectors on the drive?

No. All modern drives support error management, which completely hides any bad sectors that may be on the disk before leaving the factory. Even a single bad sector is sufficient grounds to return the drive under warranty. If you want to continue using it, the drive should be viewed with the utmost suspicion.

Western Digital's wdat_ide.exe utility can hide grown bad sectors on many Caviar disks.

There is one exception. Under rare circumstances, use of bad (too fast) timing by the disk adapter can cause bad sectors on a disk. This type of error can be fixed simply by writing fresh data to these sectors, as there is no actual media defect.

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

8.8. I keep getting CRC errors and serial overruns.

Often, this is caused by the use of block mode (see Q10.6 for an explanation). Large blocks can take a long time to transfer; during the transfer, interrupts are disabled and the serial ports are not serviced by the CPU. Eventually, the buffer for incoming data may overflow, leading to overruns and CRC errors.

The solution is to reduce the number of sectors per block, if possible, or disabling block mode altogether. 16550 compatible serial ports have a larger buffer, but with excessively large block sizes this problem may still occur.

8.9. Can my drive do 32-bit access?

There appears to be an awful lot of confusion about this subject, partly due to some unhappy terminology.

In the most literal sense, no ATA(-2,-PI) drive will allow 32-bit access. Data is transferred to and from the drive over a 16 bit bus. However, many local bus interfaces are capable of combining two 16-bit words into a 32-bit doubleword when reading data from the disk, and the reverse when writing. This way, data transfer between the CPU and the interface can be done in 32-bit chunks. This is often called '32-bit access', although '32-bit host bus transfers' would be a better name.

With 32-bit host bus transfers, more efficient use is made of the computer's bus and CPU. On the other hand, these are seldom the bottleneck, so don't expect miracles from this feature. Windows' 32-bit disk and file access are completely unrelated issues and the subject of question 8.10 and 8.11.

8.10. Help! Windows 3.x 32-bit disk access doesn't work anymore!

There are numerous reasons why this can fail; you will more easily be

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

able to do something about it (or decide if you want to fix it in the first place) once you know some background.

Windows' 32-bit disk access (32BDA) is a bit of a misnomer, actually, since it has nothing to do with 32-bit data transfers. A slightly better name for it is 'FastDisk'. It is a feature of Windows in 386 Enhanced mode that allows one to replace the BIOS' disk routines by Windows' own routines that work in protected mode. A much better name, then, would be "protected mode controller access". For some reason Microsoft decided not to use the latter.

Anyway, the main advantage of this feature is that it allows Windows to use virtual memory for its DOS sessions. Without 32-bit disk access, DOS sessions cannot be swapped out and every DOS box takes 640k of real memory. Because it also reduces the number of switches between virtual and protected mode Windows has to make, it gives a slight performance improvement as well, but usually nothing dramatic. Only if 32BDA is used together with Windows for Workgroups' 32-bit file access feature, it will eliminate these mode switches altogether (at least for most disk operations), which gives a far more interesting performance boost.

Unfortunately, the standard FastDisk routines that are internal to windows, called *wdctrl, are severely limited in their capabilities. The *wdctrl software understands nothing of non-IDE hardware (e.g. SCSI), more than two harddrives, drives with more than 1024 cylinders, 32-bit host bus transfers, block transfers, or ATAPI CD-ROM drives on the primary channel. If you use any of these things, 32-bit disk access won't work unless you have a *wdctrl replacement.

Today, that means that 32-bit disk access won't work 'out of the box' for most of us.

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

Most interfaces that are incompatible with *wdctrl come with their own FastDisk routines (usually with a .386 extension). For the rest of you, many drive manufacturers offer replacement FastDisk software. Many drive manufacturers have such drivers on their WWW sites these days; take a look in the net.resource guide below. You can also contact your vendor to find out what is available. Last but not least, the ontrackw.386 driver in [<ftp://ftp.ontrack.com/pub/software/>](ftp://ftp.ontrack.com/pub/software/) is reported to work fine on all drives even if you don't use Disk Manager.

Most of these drivers won't give you 32-bit disk access if you have an ATAPI CD-ROM on the same cable as the harddisk. Only a few CD-ROMs come with a special VxD driver which does the job.

Note: these drivers are incompatible with the Stealth feature of some versions of Quarterdeck's QEMM. Quarterdeck's fix can be found on [<ftp://ftp.wdc.com/drivers/hdutil/>](ftp://ftp.wdc.com/drivers/hdutil/).

8.11. Help! Windows for Workgroups' 32-bit file access fails!

The idiosyncrasies of the 32-bit disk access feature with respect to disk hardware has led to the popular myth that 32-bit file access has similar problems. However, that's all it is: a myth. If 32-bit file access fails, you should first check your filesystem and the programs that use it. As little as a single open file, e.g. from a printer spooler, will cause 32BFA to fail. Oh, and put `DEVICE=C:\WINDOWS\IFSHLP.SYS` in your CONFIG.SYS, and make sure your SYSTEM.INI contains the correct magic incantations (vfat.386, vcache.386). If this doesn't help, there's a first rate FAQ on this topic (see the net.resource guide for details).

8.12. Help! Win95 indicates my drive uses compatibility mode!

The culprit usually is a virus. Do get a recent virus scanner.

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

If that turns out negative, it may also be DOS (real-mode) driver that loads in the CONFIG.SYS or AUTOEXEC.BAT, or an old version of EZDrive/Disk Manager loading from the MBR.

8.13. My partitions become a mess in MS-DOS mode.

See the next question.

8.14. Win95 sees my partitions, but DOS doesn't.

If you've used Win95's fdisk utility to partition your drive, you may run across a nasty bug.

Win95 supports extended int13 calls to break the 8GB barrier. To avoid problems with old versions of DOS, partitions extending beyond 8GB must be made invisible. Unfortunately, the Win95 FDISK sometimes hides partitions this way even if your drive is much smaller than 8GB. Incidentally, this also hides them from all other operating systems, including old versions of DOS, and can cause all kinds of problems.

Under circumstances, these new partition types can completely mess up things when going from the Win95 graphical shell to MS-DOS mode. Drive contents may appear to be corrupted or be replaced by the contents of C:. Don't try anything fancy when this happens; it is really easy to corrupt your data. Don't use the "Restart in MS-DOS mode" option and don't run programs configured to run in MS-DOS mode. MS-DOS windows are still fine.

The most comfortable way to fix this is to change the partition types using Partition Magic <<http://www.powerquest.com>>, but ONLY version 2.03 or later. You can get an update patch for older versions.

The alternative is to back up your data and repartition using FDISK /X, which disables the use of the new partition types, or DOS 6 FDISK. Also be sure to apply the Win95 ios bugfix and other fixes available

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

from Microsoft's web site.

8.15. Win95 keeps losing my second IDE channel!

If you have a Triton II or Natoma based board, the retail version of Win95 may not recognize the PIIX3 interface. This will trigger an entertaining bit of Plug'n'Pray magic which eventually causes the BIOS to disable the secondary IDE channel on the next reboot.

To determine if this is really your problem, go into the device manager and click on Hard Disk Controllers. If you see the following devices listed:

- o Primary IDE Controller (single FIFO)
- o Standard Dual PCI IDE Controller
- o Standard IDEESDI Hard Disk Controller/

your Win95 mshdc.inf needs a little update. You can download this from <<ftp://ftp.intel.com/pub/patch/>>.

8.16. The Win95 busmastering drivers won't work.

The Win95 busmastering drivers sometimes have trouble co-operating with older harddisks and ATAPI CD-ROMs. Try installing the latest drivers.

If that doesn't help, you could try this registry hack. Move all old devices to the secondary port. Back up the registry (system.dat and user.dat in the Win95 directory). Start regedit and look for

```
HKEY_LOCAL_MACHINE/System/CurrentControlSet/control/Services/hdc
```

Here is where the entries for both ports should be located. In the second entry, change the key PortDriver from "ideatapi.mpd" to "esdi_506.pdr". This will cause the secondary channel to be handled by the default driver.

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

8.17. My ATAPI (IDE) CD-ROM driver won't recognize the CD-ROM.

If the CD-ROM is connected to the secondary channel, make sure this channel is enabled. Some BIOSes will enable the channel only if one or more harddisks using this channel are defined in the setup; in that case, you can't avoid putting the CD on the same cable as a harddisk until you manage to get your BIOS updated.

You may also get trouble if the CD-ROM is jumpered as slave and there's no master on its channel.

Finally, the PIO mode (speed) used by the interface may be too high, especially if the CD-ROM shares its cable with a harddisk. Many interface drivers and BIOSes are not ATAPI-aware and don't take the CD-ROM into account when determining the maximum possible speed. The best fix is to move the CD-ROM to a different channel. Manually lowering the mode a notch or two should also help; this is usually done either through the BIOS setup or by passing options to a device driver in the CONFIG.SYS.

9. The user's net.resource guide

There are a large number of FAQs and other resources on the net.

Daniel Tauritz (dtauritz@WI.LeidenUniv.NL) no longer maintains his EIDE Mini-FAQ that contains a short but potentially useful list of available interfaces; it needs to be expanded and maintained---any takers? There is another, technically oriented, FAQ on ATA-2, EIDE ! etcetera by Hale Landis (hlandis@ibm.net). All of this and more can ! be found on

- o <<ftp://ftp.wi.leidenuniv.nl/pub/faqs>>

- o <<http://www.wi.leidenuniv.nl/ata>>

An incredible range of hardware related issues is addressed by the comp.sys.ibm.pc.hardware FAQ, maintained by Ralph Valentino (ralf@wpi.wpi.edu). It can be found on comp.answers and the FAQ repository <<ftp://rtfm.mit.edu>>. It incorporates the world-famous

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

'Jumper FAQ' <<ftp://rtfm.mit.edu/pub/usenet/news.answers/faq/part3>>.

On the tangential subject of Windows for Workgroups' 32-bit file access, Mike McCormick (m.mccormick2@genie.geis.com) has a good FAQ, posted on comp.os.ms-windows.windows.setup and many other groups. It can also be found on comp.answers and therefore on rtfm.mit.edu as well.

Roedy Green's (roedy@bix.com) "PCI EIDE controller flaws" FAQ is included with his EIDE test <<ftp://garbo.uwasa.fi/pc/diskutil/>> program. The latter will test your system for the bugs that may arise if it has a CMD640 or RZ1000 interface chip (used on many pre-Triton Intel Pentium boards, for instance). With either of these, your data is in danger every day you postpone reading this.

The SimTel repository, the primary mirror of which is <<ftp://ftp.coast.net/>> or <<http://www.coast.net>>, contains an immense number of useful programs (PD, freeware and shareware). A tiny selection, all from /SimTel/msdos/diskutil/:

- o ideinf10.zip (will determine the properties of your harddisk(s)--not a benchmark).
- o 3drvs260.zip (gives support for 3 harddisks under DOS).
- o dqwik211.zip (block mode driver).
- o no_idle.zip (disables the auto-idle power saving feature of some harddrives).
- o lseagate.zip (specs for *all* Seagate harddrives).
- o wasted15.zip (shows wasted space due to cluster size).
- o presz111.zip (nondestructive partition resizer). r

Note that the version numbers may have changed; also, you should preferably use a mirror close to you instead of the oakland site.

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

A more manageable selection of useful utilities, drivers and information (this FAQ :-) can be found on

- o [<ftp://ftp.netcom.com/pub/cl/clau/>](ftp://ftp.netcom.com/pub/cl/clau/)

(this represents the latest known location of Chung Y. Lau's famous wandering FTP site)

- o [<ftp://ftp.rahul.net/pub/lps/>](ftp://ftp.rahul.net/pub/lps/)

One of the most thorough low level harddisk benchmarks on Earth, Marnix Timmermans' Check Harddisk, is in beta. Check it out at [<http://huizen.dds.nl/~checkhd>](http://huizen.dds.nl/~checkhd). Of course, the usual precautions with beta software apply.

SimTel has a copy of TheRef[TM] by F. Robert Falbo, a giant hardware reference; it is rather outdated, but fortunately you can access a more up to date copy on the web at [<http://theref.c3d.rl.af.mil>](http://theref.c3d.rl.af.mil).

Frank Pikelner (frank@cs.yorku.ca) has compiled an excellent, up to date list of >500MB harddisk and >4x CD-ROM drive specifications. You can access it through

[<http://www.cs.yorku.ca/People/frank/Welcome.html>](http://www.cs.yorku.ca/People/frank/Welcome.html).

9.1. The user's net.resource guide: I/O card drivers

! DriversHQ [<http://www.drivershq.com>](http://www.drivershq.com) is one of the most well known
! sources for drivers.

!

In Finland, someone got the immensely useful idea of setting up a site with all kinds of IDE card drivers. You can find it on

[<ftp://ftp.funet.fi/pub/drivers/pc/disk/ide_ata>](ftp://ftp.funet.fi/pub/drivers/pc/disk/ide_ata). If you can't find a driver here and manage to pick it up somewhere else, please upload it. Help making the net work.

9.2. The user's net.resource guide: hardware manufacturers

A number of disk and controller manufacturers now have FTP and web

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

sites; these are

- o CMD

<<http://www.cmd.com>>

- o Conner

<<http://www.conner.com>>

- o Creative Labs

<<http://www.creaf.com>>

Among other, non EIDE related items, this site has drivers for the tertiary ATA/IDE port found on some SoundBlaster cards.

- o DTC

<<http://www.datatechnology.com>>

- o Fujitsu

<<http://www.fujitsu.com>>

- o Hewlett-Packard

<<http://www.hp.com>>

- o IBM

<<http://www.storage.ibm.com/storage>>

<<ftp://hddtech.millcomm.com>>

<<http://www.pc.ibm.com>> (PC Company)

- o Iomega

<<http://www.iomega.com>>

- o Maxtor

<<http://www.maxtor.com>>

<<ftp://ftp.maxtor.com>>

All files from their BBS are on the FTP site, including biosbnch.zip, a benchmarking utility; greendrv.zip, to set the sleep mode timer on 'green' drives; and the latest version of MaxBlast (OnTrack Disk Mananager) at <<http://www.maxtor.com/readme/dm.html>>.

- o Promise

<<http://www.promise.com>>

<<ftp://ftp.promise.com>>

support@promise.com

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

This site doesn't only contain drivers for Promise products but also for clone cards such as the VG4.

- o Quantum

[<http://www.quantum.com>](http://www.quantum.com)

No official FTP site, although drivers to use Quantum's large harddisks with a non-translating BIOS can be found in [<ftp://ftp.rahul.net/pub/lps/hard-disk/>](ftp://ftp.rahul.net/pub/lps/hard-disk/). This is NOT an official Quantum support site and may disappear. This site contains miscellaneous storage-related programs and information too, not least the white papers on topics such as S.M.A.R.T., PRML channels, MR heads and more: [<http://www.quantum.com/products/whitepapers/>](http://www.quantum.com/products/whitepapers/).

- o Seagate

[<ftp://ftp.seagate.com>](ftp://ftp.seagate.com)

[<http://www.seagate.com>](http://www.seagate.com)

This includes specifications of ALL Seagate drives, including detailed diagrams/pictures. The Seagate 32-bit disk access driver is available as [;<ftp://ftp/seagate.com/techsuppt/sea32bit.exe>](ftp://ftp/seagate.com/techsuppt/sea32bit.exe); there is a FAQ on this driver at [<http://www.seagate.com/techsuppt/win32drv.html>](http://www.seagate.com/techsuppt/win32drv.html). There is some general technical information available, on topics such as MR heads: [<http://www.seagate.com/new/sep96/mr_techp.shtml>](http://www.seagate.com/new/sep96/mr_techp.shtml). A formatting utility for Seagate drives, `sgatfmt4`, is also available.

- o Tekram

[<http://www.tekram.com>](http://www.tekram.com)

- o Tyan

[<http://www.tyan.com>](http://www.tyan.com)

- o Western Digital

[<ftp://ftp.wdc.com>](ftp://ftp.wdc.com)

[<http://www.wdc.com>](http://www.wdc.com)

Many drivers and utilities, some of which won't work if you don't have at least one WD Caviar drive in your system. Goodies to be found: a good Windows FastDisk driver in `/drivers/hdutil/win31.exe`; a Windows helpfile explaining Enhanced IDE in `/docs/eide.exe`; a

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

utility to examine the Enhanced DPT of your BIOS:

`/drivers/hdutil/chkbios.com`.

- o Winbond

`<http://ntwww.winbond.com.tw>`

Drivers are in the `/DNLOAD/` directory.

Pointers to sites offering software from other manufacturers would be appreciated.

9.3. The user's net.resource guide: software and BIOS houses

- o AMI

`<http://www.megatrends.com>`

`<ftp://ftp.megatrends.com>`

- o Award

`<http://www.award.com>`

- o IBM

`<http://www.ibm.com>`

`<http://www.pc.ibm.com>` (PC Company)

`<http://ps.boulder.ibm.com>` (a really good link for OS/2 install and update info)

`<ftp://ftp.pc.ibm.com>`

Flash BIOS upgrades for the IBM ValuePoints are in `/pub/valuepnt/`.

- o MicroHouse

`<http://www.microhouse.com>`

Drive specifications and jumper settings:

`<http://www.microhouse.com/mtl/vip/hd.htm>`. EZDrive FAQs and files can be found on `FAQ/ALLKEY.htm` and `FTP/EZ.htm` respectively.

- o Microsoft

`<ftp://ftp.microsoft.com>`

`<http://www.microsoft.com>`

General Windows and DOS updates, fixes and Microsoft Knowledge Base articles can be found on the Microsoft FTP site. Unfortunately, the structure of this site can hardly be called intuitive (which, some would argue, is characteristic of Microsoft products in

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

general). The WWW site allows searches in the Knowledge Base, which is much more useful.

- o MR BIOS

<<http://www.mrbios.com>>

mrbios@mrbios.com

Information on MR BIOS bioses. Contains downloadable shareware BIOS images for some popular boards with Flash ROM as well.

- o Novell

<<http://www.novell.com>>

<<http://www.netware.com>> (Netware info)

- o OnTrack

<<http://www.ontrack.com>>

tech@ontrack.com (tech support)

sales@ontrack.com (sales)

- o Phoenix

<<http://www.ptltd.com>>

- o SCO

<<http://www.sco.com>>

- o Unicore software

<<http://www.unicore.com>>

Calls itself a BIOS solutions site. Contains, among other things, information about the LBA Pro BIOS add on board.

10. ATA: harddisks

This and the following sections of the FAQ are intended to provide additional background information for the curious. The answers here differ wildly in the depth in which they treat the material; pick and read at will. It's a FAQ, not a book.

10.1. How does ATA(-2) work?

To understand the basic concepts of advanced ATA drives, one first needs to understand the basics of drive technology. Basically, when the operating system needs data to be either read or written to

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

secondary storage (the hard disk), the BIOS gets the command, and passes that command to the drive. For operating systems other than DOS, the BIOS is usually replaced by the operating system's own I/O subsystem; the principle remains the same.

How the command is passed, interpreted, and responded to, forms the basis for Advanced ATA. In a nutshell, there are seven registers that the BIOS writes to/reads from to create a command. An eighth register is used to read and write data. The signals that create these reads and writes are controlled by the BIOS, but their timing is determined by the interface hardware, and ATA specifications dictate how fast these signals can be asserted or deasserted. There are currently 4 modes of Programmed Input/Output (PIO) and 4 modes of Direct Memory Access (DMA). The numbers all of you have been reading about are only a small portion of these specifications, but they are the ones that marketing can tout best. These "transfer rates" are a result of the specification that controls how fast the I/O Read and Write cycle time of the data register can operate at.

10.2. What are PIO modes?

The PIO mode determines how fast data is transferred to and from the drive. In the slowest possible mode, PIO mode 0, the data cycle time can not exceed 600 nanoseconds. In a single cycle, 16 bits are transferred in or out of the drive. In a single sector, there are 256 words (16 bits = 1 word); 2048 sectors make up a megabyte. So, mathematically,

$$\begin{array}{rcccccc} 1 \text{ cycle} & 1 \text{ sector} & 1 \text{ megabyte} & & 2000 & \\ \text{-----} & \text{-----} & \text{-----} & = & \text{-----} & = 3.3\text{MB/s} \\ 600\text{ns} & 256 \text{ words} & 2048 \text{ sectors} & & 600\text{ns} & \end{array}$$

So, the theoretical transfer rate of PIO Mode 0 (600ns cycle time) is 3.3 megabytes per second.

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

Here are the rest of the PIO modes, with their respective transfer rates:

| PIO mode | Cycle time (ns) | transfer rate (MB/s) | |
|----------|--------------------|-------------------------|-----------------------|
| 0 | 600 | 3.3 | ATA |
| 1 | 383 | 5.2 | ATA |
| 2 | 240 | 8.3 | ATA |
| 3 | 180 | 11.1 | ATA-2, IORDY required |
| 4 | 120 | 16.6 | ATA-2, IORDY required |
| 5 | 90 | 22.2 | vaporware |

The first three, PIO modes 0 to 2, are old modes also present in the old ATA standard. The others (PIO 3 and 4) are ATA-2 specific and use IORDY hardware flow control. This means the drive can use the IORDY line to slow down the interface when necessary. Interfaces without proper IORDY support may cause data corruption in the fast PIO modes; in that you're stuck with the slower modes, and typically half the bandwidth.

When interrogated with an Identify Drive command, a harddisk returns, among other things, information about the PIO and DMA modes it is capable of using.

10.3. What are DMA modes?

DMA or Direct Memory Access means that the data is transferred directly between drive and memory without using the CPU as an intermediary, in contrast to PIO. In true multitasking operating systems like OS/2 or Linux, DMA leaves the CPU free to do something useful during disk transfers. In a DOS/Windows environment the CPU will have to wait for the transfer to finish anyway, so in these cases DMA isn't terribly useful.

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

There are two distinct types of direct memory access: third-party DMA and first-party or busmastering DMA. Third-party DMA relies on the DMA controller on the system's mainboard to perform the complex task of arbitration, grabbing the system bus and transferring the data. In the case of first-party DMA, all this is done by logic on the interface card itself. Of course, this adds considerably to the complexity and the price of a busmastering interface.

Unfortunately, the DMA controller on ISA systems is ancient and slow, and out of the question for use with a modern harddisk. VLB cards cannot be used as DMA targets at all and can only do busmastering DMA. It is only on EISA- and PCI-based interfaces that non-busmastering DMA is viable: EISA type 'B' DMA will transfer 4MB/s, PCI type 'F' DMA between 6 and 8MB/s.

Today, all modern chipsets, including the ubiquitous Triton chipsets, incorporate a busmastering DMA capable ATA interface. Efforts to standardize the DMA hardware will ensure stable and reliable software support.

Anyway, the DMA modes supported are:

| | DMA Mode | Cycle time | transfer rate | |
|---|-------------|------------|---------------|-----------|
| | Single word | (ns) | (MB/s) | |
| | 0 | 960 | 2.1 | ATA |
| | 1 | 480 | 4.2 | ATA |
| | 2 | 240 | 8.3 | ATA |
| | Multiword | | | |
| | 0 | 480 | 4.2 | ATA |
| | 1 | 150 | 13.3 | ATA-2 |
| | 2 | 120 | 16.6 | ATA-2 |
| ! | DMA/16 | 120 | 16.6 | Ultra-ATA |
| ! | DMA/33 | 60 | 33.3 | Ultra-ATA |

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

The single word DMA modes are hardly useful and are obsoleted in ATA-3. Note that some older interfaces are able to use these DMA modes as a way to communicate with the drive, without actually doing direct memory access at all. In these cases, the DMA modes are just used as glorified PIO modes.

10.4. How are the ATA(-2,PI) I/O ports assigned?

The registers of the primary ATA channel occupy the following I/O addresses (in hexadecimal notation):

| Register | Read Function | Write Function |
|----------|--------------------------------|---|
| 01F0h | Read Data (16 Bits) | Write Data (16 bits) |
| 01F1h | Error register | Set Features Data |
| 01F2h | Status of sector count | Write sector count for command setup |
| 01F3h | Location of starting sector | Write sector start for command setup |
| 01F4h | Location of Cyl-low | Write cyl-low location for command setup |
| 01F5h | Location of Cyl-high | Write cyl-high location for command setup |
| 01F6h | Head/device selection | Write device selection and head selection for command setup |
| 01F7h | Device Status | Device command |
| 03F6h | Alternate Status | Device Control |
| 03F7h | Drive Address | |

Note that the floppy disk controller's disk change flag shares 03F7h which makes life difficult for designers that want to implement disk and floppy controllers separately. From this point of view it may come

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

as no surprise that some of the problems in the CMD640x and RZ1000 'EIDE' interface chips touched upon elsewhere in this FAQ are floppy related.


There is no reason why there can't be a large number of interfaces like this one. There is a de facto standard for four of these ports:

| Interface number | CS0-decode | CS1-decode | IRQ number |
|------------------|-------------|-------------|------------|
| 1 | 01F0h-01F7h | 03F6h-03F7h | 14 |
| 2 | 0170h-0177h | 0376h-0377h | 15 or 10 |
| 3 | 01E8h-01EFh | 03EEh-03EFh | 12 or 11 |
| 4 | 0168h-016Fh | 036Eh-036Fh | 10 or 9 |

Only the first two enjoy really widespread support; for the secondary port, IRQ15 is the most commonly used interrupt by far. Potential BIOS support for arbitrary extra ports is found only in the Phoenix specification.

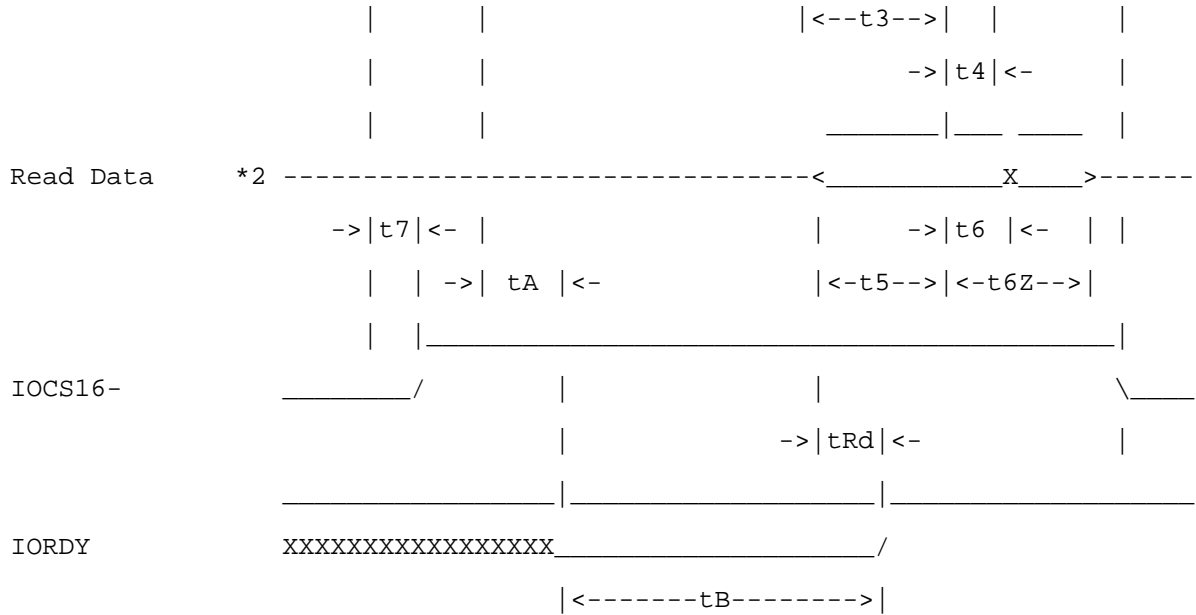
10.5. What does an ATA-2 interface do?

Interfaces have come a long way since the ordinary ISA IDE consisting of little more than a simple buffer. ATA-2 boards have to support at least PIO modes 0 and 3, usually support many more modes, and will have to ensure that the correct timing is used at the ATA interface for each of these modes. Since the timing specifications are quite complicated, a great deal of flexibility is necessary to implement the ATA-2 standard correctly.



Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan



*1 Device Address consists of signals CS0-, CS1- and DA2-0

*2 Data consists of DD0-15 (16-bit) or DD0-7 (8-bit)

The above figure defines the relationships between the interface signals for both 8-bit and 16-bit PIO data transfers.

In this diagram, t_0 denotes the read/write cycle time, the most significant determining parameter for PIO mode throughput. As you can see, there is a lot more to the various PIO and DMA modes than this read/write cycle time only. To design a low-cost interface that fully adheres to the ATA-2 specification is quite a challenge. The common approach is to make the timing completely software programmable; unfortunately, the way ATA cards are programmed has not been standardized and differs radically between cards. The consequence is that you will typically need interface-specific drivers for each and every operating system used in order to profit from the fast transfer modes.

10.6. What is Block mode?

Multiple Read/Write commands (reduces Interrupts to host processor). Besides the obvious transfer increase, Fast-ATA and many other drives allow for Read/Write Multiple commands, which increase the number of

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

sectors passed without intervening interrupts. This lessens the host's overhead, as every interrupt causes the CPU to do a context switch, check the device and set up the data transfer (or perform the transfer itself in the case of PIO).

The Read Multiple Command (0C4h) and the Write Multiple Command (0C5h) are drive-level commands that can transfer multiple sectors of data without asserting the IRQ line of the drive, signaling the processor that a drive operation is pending.

The IRQ line is asserted when:

- o A read command has been issued, and the requested data is in the drive's buffer, ready to be taken by the host.
- o A write command has been issued, and the data has transferred to the drive's buffer. If write caching is disabled, the IRQ won't be asserted until the data has been completely written to the media.

During normal reads and writes, the interrupt can constantly bother the CPU, and depending on the processor and the task at hand (multi-tasking OS, Unix, etc), there can be long delays in having the CPU service the drive. The advent of Read/Write Multiple allows many sectors (from 2 up to as many as 128) to be transferred in one go, completing the task in as much as 30% faster times.

On single-tasking operating systems like DOS, any improvement over a few percent usually indicates bad buffer cache management on the part of the drive.

Warning: some old drives have a buggy block mode implementation and may corrupt data.

A final remark: the block size that is optimal for drive throughput doesn't have to be the best for system performance! For example, the DOS FAT filesystem tends to favor a block size equal to the cluster size. Do not trust low level benchmarks when tweaking the block size,

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

but use an application level benchmark suite instead.

10.7. What is LBA?

LBA is a means of linearly addressing sectors addresses, beginning at sector 1 of head 0, cylinder 0 as LBA 0, and proceeding on to the last physical sector on the drive, which, for instance, on a standard 540 Meg drive would be LBA 1,065,456. This is new in ATA-2, but has always been the one and only addressing mode in SCSI. Note that LBA does not allow you to address more sectors than CHS style addressing would.

LBA reduces CPU overhead in OSs that use LBA internally, but on the other hand takes a little more time when ordinary CHS based BIOS calls are used (eg. DOS). Beware that depending on the way LBA is implemented in the harddisk firmware, the overhead on the part of the drive may increase.

10.8. How does security work?

Security mode implements a simple password protection scheme. Security can affect just write operations, or both reads and writes. Non-data operations (such as Identify Device) can always be executed regardless of the security status.

There are two security levels: High and Maximum. If the user password is lost and High level security is set, the drive can still be unlocked with the Master password. At Maximum level, there is no way to unlock the drive without erasing all user data.

After a number of incorrect passwords the drive will reject further passwords until a powerdown or hard reset. This makes guessing the password by brute force very difficult.

10.9. What is S.M.A.R.T.?

S.M.A.R.T. or Self Monitoring Analysis and Reporting Technology allows

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

the drive to report about certain types of degradation or impending failure. This allows the operating system to take the necessary precautions and warn the user. The OEM release 2 of Win95 and the next OS/2 version (Merlin) will be SMART aware.

The utility of this feature will initially be quite limited, though, because many failure modes (including the infamous Monday morning failure) can't be sensed in advance.

At present only few utilities exist to examine the S.M.A.R.T. status of a drive. These include Micro House EZ-S.M.A.R.T. and Symantec S.M.A.R.T. Doctor.

10.10. What is PRML?

The Partial Response Maximum Likelihood or PRML read channel is quickly replacing the ordinary peak detection channel as a mass market technology. Briefly, where a peak detection channel uses a comparatively simple analogue technique to extract the digital data from the signal picked up by the read head, a PRML channel digitizes the signal and employs digital processing techniques to reconstruct the data. Thanks to these DSP techniques a PRML channel can still work reliably on very closely packed data where the distinction between individual bits tends to blur. The end result is a faster, higher capacity drive.

For a more thorough discussion of this topic, take a look at Quantum's excellent white papers <<http://www.quantum.com/products/whitepapers/>>.

10.11. What are MR heads?

Magnetoresistive or MR drive heads are quickly replacing the old inductive heads in all sections of the harddrive market. For more information, see for example <http://www.seagate.com/new/sep96/mr_techp.shtml>.

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

11. ATAPI: CD-ROMs and tapes

11.1. How does ATAPI differ from, and coexist with, ATA(-2)?

For the sake of compatibility with non-ATAPI aware software that might mistake an ATAPI device for a harddrive, the device pretends it isn't there until it's waken up by a special sequence of commands. Once activated, it uses a command protocol that radically differs from that used by harddisks.

The reason is that the ATA command and register set is not adequate to support some CD-ROM command structures. Therefore, only a minimum of traditional ATA commands are supported by ATAPI devices. For most of their functions, these devices rely on the ATAPI Transport Protocol using packets of at least 12 bytes sent, as data, through the Data Register. These packet commands have been derived from the SCSI command set; this makes it reasonably easy to rewrite existing SCSI CD-ROM and tape drivers for ATAPI hardware.

Beware that non-ATAPI aware 'intelligent' controllers, mainly caching controllers, will be mightily confused by packet commands.

Traditionally, the data register is only used to transport 512-byte sectors; a 12-byte command packet is a completely different kind of animal and should be treated in a different way by the (intelligent) controller.

11.2. What's so special about the secondary port?

Nothing, in principle. A secondary IDE port has been reserved in the PC I/O map for ages (base address 0170h, IRQ 15), and adapters that could be configured as secondary have been available for quite some time, even though BIOS support was lacking. So while it is a part of EIDE, there is actually nothing new about this feature, except that the possibility of connecting tape drives and CD-ROMs to the ATA adapter has transformed four device support from a luxury into a necessity.

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

Actually, there is another reason to provide a secondary port for ATAPI devices. There are a number of advanced hardware features for harddisk interfaces, such as prefetch buffers and write behind, that may get in the way of ATAPI compatibility. This means that if the software drivers of an intelligent ATA-2 port are not ATAPI-aware, or simply don't work as they should, you may run into sticky problems. Especially if you have an ATAPI CD-ROM that doesn't support PIO mode 3 transfers, a secondary port that provides only basic ATA features is a good way to avoid a lot of headaches.

Finally, an ATAPI device on the primary port will cause Windows FastDisk drivers that aren't ATAPI aware to fail.

Nothing prevents you from defining more ports like the primary and the secondary one; in addition to these two, the tertiary and quaternary one are semi-standard. See section 10.4 for the port and IRQ assignments of all these ports.

12. The EBIOS: translation

12.1. Why translation?

Both the 'int13' software interface used by the BIOS to communicate with the outside and the Cylinder/Head/Sector (CHS) fields in the partition table reserve

- o 10 bits for the cylinder field, for a total of up to 1024 cylinders;
- o 8 bits for the head field, good for up to 256 heads;
- o 6 bits for the sector field, which gives a maximum of 63 sectors since for historic reasons the sector field starts at sector 1, not 0.

The maximum disk capacity accessible through the traditional int13 interface is therefore 8GB (1024*256*63 sectors of 512 bytes). In some books, you may encounter references to 12-bit cylinder

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

numbers; this extension (using the upper two bits of the sector field) was never widely implemented and isn't supported anywhere.

Now IDE disks have their own set of limitations; these disks, no matter if they're ATA/IDE or ATA-2/EIDE, use

- o 16 bits for the cylinder field, giving 65536 cylinders;
- o 4 bits for the head field, or only 16 heads at most;
- o 8 bits for the sector field.

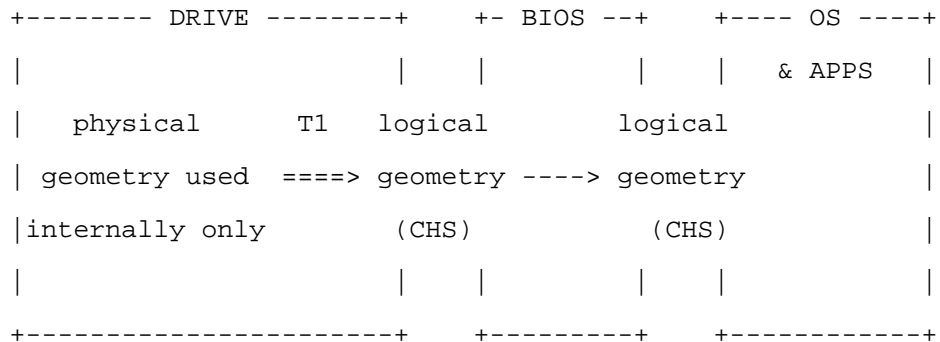
This is good for a maximum disk capacity of 128GB. However, combine this with the BIOS limitations and you suddenly can't see more than the first 1024 cylinders of the IDE disk, which makes for a limit of just 504MB or 528 million bytes. This is unacceptable today. In the long term, the BIOS limit of 8GB is just as unacceptable, but as a short term solution it is desirable to get the maximum out of the standard int13 interface with IDE drives. This is where translation comes in.

12.2. How does translation work?

There are roughly three ways today's BIOSes can handle translation: standard CHS addressing, Extended CHS addressing, and LBA addressing. Translation does NOT automatically imply LBA, as you will see.

- o Standard CHS: no translation at all :-(

Communication between drive, BIOS and operating system goes like this:



There is only one translation step, T1, which is internal to the drive ('universal translation'). The drive's actual, physical geometry is completely invisible from the outside---the Cylinders,

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

Heads and Sectors printed on the label for use in the BIOS setup have nothing to do with the physical geometry! The logical geometry, used throughout, is subject to both IDE's limitation to 16 heads and to the BIOS' limitation of 1024 cylinders, which gives the (in)famous 504MB limitation.

o Extended CHS

```
+----- DRIVE -----+   +- BIOS ---+   +----- OS -----+
|           |           |   |           |   |           |   & APPS   |
|  physical   T1  logical  T2  translated           |
| geometry used  =====> geometry =====> geometry           |
|internally only      (CHS)                (CHS)           |
|           |           |   |           |   |           |   |
+-----+-----+   +-----+   +-----+
```

Logical geometry is used to communicate between the drive and the BIOS, while a different, translated geometry is used to communicate between the BIOS and everything else.

There is an additional translation step, T2, performed by the BIOS. This procedure breaks the 504MB barrier because the geometries used are not subjected to the BIOS and IDE limitations simultaneously: the logical geometry is subject to IDE's 16 head limitation, but not to the 1024 cylinder limitation. For the translated geometry, it is just the reverse.

Most BIOSes denote extended CHS translation with 'Large'. Note that the geometry usually entered in the BIOS setup is the logical geometry, not the translated one. In case of doubt, consult the BIOS manual.

o Logical Block Addressing (LBA)

Here, the logical geometry is dispensed with entirely and replaced by a single, large, linear block number. This makes far more sense than using a logical geometry that is completely fake anyway.

```
+----- DRIVE -----+   +- BIOS ---+   +----- OS -----+
|           |           |   |           |   |           |   & APPS   |
```

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

| | | | | | | |
|--|-----------------|-------|----------------|----------|------------|--|
| | physical | T1 | linear | T2 | translated | |
| | geometry used | ====> | block no.====> | geometry | | |
| | internally only | | (LBA) | | (CHS) | |
| | | | | | | |
| | +-----+ | | +-----+ | | +-----+ | |

This breaks the 504MB barrier in essentially the same way as extended CHS does. Conceptually, using a single linear number to address a sector on the harddisk is simpler than a CHS style address, but it takes more CPU cycles and is sometimes slower on the drive side as well. The differences are pretty insignificant either way.

A translating BIOS can be implemented via the system BIOS or on-board controller BIOS. Basically, this takes the drive's logical default geometry, and if the cylinder count is beyond 1024, will divide the cylinder count by an appropriate factor and multiply the heads by the same. For instance, let's take a 540 Meg drive: it has 1057 cylinders, 16 heads, and 63 sectors per track. Well, the int13 interface used by the BIOS to talk with the world can only handle 1024 cylinders, but it can address up to 255 heads. So, the x-
lating BIOS will pass to the OS, via int13 calls, the geometry 528 cylinders (1057/2 (approx)) and 32 heads (16*2). Then, when the OS makes a request to the drive, the BIOS will re-translate the request to the original order, or to an LBA number if LBA is enabled. This allows for capacities of up to 8 gigabytes.

A final word about the 8GB capacity limit, which is inherent in the int13 interface and cannot be solved without ditching the traditional calls. To that purpose, the IBM/Microsoft int13 extensions document specifies a new interface between the BIOS and the operating system or applications. These extended int13 calls are made in terms of LBA addresses and can handle huge disks. Note that the BIOS is required to translate these LBA addresses back to CHS if the drive doesn't support

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

LBA---exactly the reverse of the translation process outlined above.

12.3. I'd like to know how translation works in detail.

You asked for it :-)

If a drive is less than 504MB, it should have a logical geometry, as reported in Identify Device words 53-58, of 1024 or less cylinders, 16 or less heads and 63 or less sectors. Such a drive can be addressed directly without invoking this algorithm. For drives over 504MB, the CHS address received by the BIOS from DOS must be converted to an Extended CHS address, or an LBA address. We'll assume ECHS for now. First, during BIOS setup, the BIOS must determine the value of N. This value is used to convert the drive's geometry to a geometry that the BIOS can support at the int13 interface. This interface requires that Cyl be less than or equal to 1024. The number of cylinders (Identify Device word 1) is divided by N while the number of heads (Identify Device word 3) is multiplied by N. N must be 2, 4, 8, ..., a power of 2.

Second, in most translating BIOSes, the following algorithm is used whenever INT 13H is called to perform a read or write:

```
eCyl    = ( Cyl * N) + ( Head / dHead ); /* head DIV dHead */
eHead   = ( Head % dHead );             /* head MOD dHead */
eSector = Sector;                       /* used as is      */
```

By way of example, assume the drive's geometry is 2000 cylinders, 16 heads and 63 sectors (these numbers are in Identify Words 1, 3, and 6) and that the BIOS determines the value of N to be 2. The BIOS reports to DOS that the drive has 1000 cylinders, 32 heads and 63 sectors when int13 ah=08h function is called. This is 2016000 sectors.

| Cyl/Head/Sector | eCyl/eHead/eSector | LBA |
|-----------------|--------------------|---------|
| 0/0/1 | 0/0/1 | 0 |
| : | : | : |
| 500/0/1 | 1000/0/1 | 1008000 |

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

| | | |
|-----------|------------|---------|
| 500/15/63 | 1000/15/63 | 1009007 |
| 500/16/1 | 1001/0/1 | 1009008 |
| 500/31/63 | 1001/15/63 | 1010015 |
| 501/0/1 | 1002/0/1 | 1010016 |
| : | : | : |
| 999/31/63 | 1999/15/63 | 2015999 |

Note the following about this algorithm: The physical ordering of sectors on the drive is unaffected---sector n is followed by sector n+1 for all CHS and Extended CHS and LBA addresses. This is the only sane way of implementing translation, but unfortunately NOT ALL BIOSSES DO IT THIS WAY. This means that changing translation modes may be a dangerous thing to do.

12.4. What is in the Enhanced Disk Parameter Table?

In a standard BIOS, the Fixed Disk Parameter Table (FDPT) contains information about the geometry of the harddisk(s). It more or less contains the same information as the drive type entry in the CMOS setup. A program that wants to use the harddisk on a low level, bypassing DOS, normally uses the BIOS' int13 functions to achieve this.

The Enhanced fixed Disk Parameter Table (EDPT) is an extension of the ordinary FDPT that makes use of undefined fields to provide information about the translation mode used. It uses a magic number (A0 in byte 3) and a checksum (in byte 15) to ensure that software cannot mistake random data for an EDPT. This practice is more or less standard across various flavors of translating BIOSes, with differing magic numbers.

On top of this, the Phoenix Enhanced BIOS standard specifies a number of extended int13 functions and a 16 byte FDPT extension. The latter contains detailed information about the current PIO or DMA type, block mode used, LBA or (E)CHS addressing, 32 bit transfer mode, media type

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

(removable, CD-ROM), control port base and IRQ. In other words, it covers all new features of the ATA family and is flexible enough to accommodate more than four devices, nonstandard port addresses and IRQs. Proper support of all this is important to achieve any degree of Plug'n'Play functionality with ATA hardware. The WD EIDE BIOS lacks all these features.

12.5. How many types of translating/Enhanced BIOSes are there?

Too many. See question 12.4 for a discussion of some of the differences between a Phoenix EBIOS and a WD EBIOS. For a more exhaustive discussion of BIOS types, Hale Landis' effort is indispensable---see the resource guide below.

13. Software details

13.1. Details on OnTrack Disk Manager

Disk Manager 6.x and above is a piece of software that performs the translation necessary to access harddisks of more than 1024 cylinders with DOS/Windows. This is achieved by installing a Dynamic Drive Overlay (DDO) to translate drive parameters. The driver provides only the basic translation functions, without EDPT or extended int13 calls. Of course, this is less of an issue in a software driver than in a system BIOS.

If Disk Manager (DM) is used to format only the slave drive, the DDO can be installed in the config.sys as an ordinary device driver (device=dmdrvr.bin). On the other hand, using DM on the master drive isn't quite that easy from a technical point of view. Since you must boot DOS from the master drive, and you must load the DDO before you can access the drive, the only option is to load it very early during the boot process, even before the operating system itself. Changes are made to the Master Boot Record (MBR) to accomplish this 'pre-boot loading'.

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

This scheme works fine, but has a few drawbacks.

- o First, DDO as a pre-boot loader has implications for floppy boots. If you boot directly from a floppy, DDO does not have a chance to boot and the partition does not make sense. This can be remedied by having a line in the config.sys on the floppy which reads "device=dmdrvr.bin". You can also watch for the press spacebar to boot from floppy message when booting from the hard drive.
- o The second drawback is that operating system installations routinely overwrite the MBR with their own boot code. The DDO is no longer loaded and your partition will be inaccessible until you let Disk Manager write a new MBR.
- o Third, a nonstandard partitioning scheme is used whereby the data is offset by one track. This is completely transparent as long as the drive is accessed through the DDO; however, many operating systems want to replace the DDO by their own disk routines and will have to be aware of this scheme.

Windows 95 will support Disk Manager 6.x and above 'out of the box', as will new versions of OS/2 Warp, Windows NT 3.5.1 and Linux 1.3.x. IBM and Microsoft have created fixes that allow older versions of OS/2, Warp and Windows NT to work with Disk Manager.

- o Fourth, corruption of the DDO sector will result in a DDO Integrity Error. While it is fairly easy to re-write the DDO sector (use the dmcfig.exe utility on your DM diskette), this is a sign of a bigger problem (eg. virus) rather than a problem in itself---contact Ontrack tech support (tech@ontrack.com) for assistance.

An advantage of formatting the master drive with DM instead of loading the DDO from config.sys is that you can use Windows for Workgroups' 32-bit file access on both drives---if you use dmdrvr.bin, the slave drive is restricted to 16-bit file access.

The 6.x versions of Disk Manager have some additional disadvantages

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

which are corrected in version 7:

- o They are not fully compatible with the device drivers of most VLB ATA(-2) interfaces; also, ATAPI CD-ROM and tape devices on the chain are not supported.
- o A final concern is disk utilities. If the utility in question goes directly to the hardware, without going through the DD overlay, it can potentially be destructive. Ontrack's policy on this is to refer compatibility questions to the manufacturer of the utility as they cannot possibly maintain compatibility charts for all versions of all utilities.

You can find more information on the various versions of Disk Manager on OnTrack's www site <<http://www.ontrack.com>>.

13.2. How does Windows' 32-bit disk access work?

32-bit disk access (32BDA), also known as FastDisk, is a set of protected-mode drivers that direct int13 calls to the hard disk controller through a protected mode interface. For the latter the hard disk controller has to supply an appropriate virtual device driver (VxD).

Windows ships with one such driver built in: *wdctrl. Unfortunately, this device only supports controllers that are strictly compatible with the WD1003 standard; this excludes SCSI, ATA-2, LBA or CHS translation, disks with more than 1024 cylinders and even some commonplace features of ATA such as block mode. If it detects one of these during the initialization phase it will refuse to load. In today's computers, this means that *wdctrl will rarely do the job and an external VxD must be used.

32BDA has two advantages over disk access through the BIOS. First, since the FastDisk VxD is re-entrant, it enables Windows to use virtual memory for DOS sessions. Using virtual memory without 32BDA could create a deadlock situation if a page fault is generated during

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

the execution of BIOS routines. Since the BIOS is not re-entrant, it is not possible to use a BIOS call to read the page from disk until the first BIOS call has terminated; on the other hand, this BIOS thread must remain suspended until the swapped out page has been read.

So 32BDA enables Windows to manage memory much more efficiently with one or more DOS sessions open.

The second advantage of 32-bit disk access is that it saves two (relatively slow) switches between virtual and protected mode per disk I/O call. Take, for instance, a disk read performed by a DOS application. In the absence of 32BDA, each such call causes the following sequence of events:

- 1 Application calls INT21 to read from disk
- 2 Windows traps the call, switches to protected mode
- 3 Windows switches to real mode, returns to DOS
- 4 DOS makes int13 call to BIOS disk routines
- 5 Windows traps the call, switches to protected mode
- 6 Windows switches to real mode, returns to BIOS
- 7 BIOS acts upon int13 call and does the read
- 8 Windows traps the return from int13, switches to PM
- 9 Windows switches to RM, returns the result to DOS
- 10 DOS receives the result, passes on to application
- 11 Windows traps the return from DOS, switches to PM
- 12 Windows switches to RM, returns result to application
- 13 Application receives the result from the INT21 call

Using 32-bit disk access replaces steps 6 to 8 by a single call to the FastDisk VxD. This removes two mode switches, resulting in a usually small disk performance improvement. (Steps 3-11 also apply to native Windows applications).

14. Hacker's resource guides

14.1. The hacker's documentation guide

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

Unfortunately, this lists consistently eludes being entirely up to date, but it should get you started.

- o AT Attachment Interface for Disk Drives, ANSI X3.221-1994, Approved May 12, 1994.
- o AT Attachment Interface with Extensions (ATA-2), ANSI ASC X3.279-1996, revision 3, proposed American National Standard 948D.
- o AT Attachment-3 Interface (ATA-3), ANSI ASC X3.298-199x.
- o AT Attachment-4 Interface (ATA-4), X3T13 draft.
- o ATA packet Interface for CD-ROMs, SFF-8020, Revision 1.2, June 13 1994.
- o Western Digital Enhanced IDE Implementation Guide, by Western Digital Corporation, revision 5.0.
- o Fast ATA Sourcebook, Quantum Corporation, November 1994.
- o Enhanced Disk Drive Specification, by Phoenix Technologies Ltd., version 1.1, January 95.

14.2. The hacker's net.resource guide

! Hale Landis (hlandis@ibm.net) has written a number of "how it works" documents dealing with boot sectors, MBRs and partition tables. These are available in a ZIP archive from

<ftp://fission.dt.wdc.com/pub/otherdocs/pc_systems/how_it_works/>.

Also, he has an extensive description of BIOS types, invaluable for programmers, a "Facts and Fiction" series that dispels a couple of common myths, and more. These, together with the How It Works

documents (in case you cannot FTP) are available from his mail server; to get started, send an e-mail message

To: hiw@sugs.tware.com

help

end

This server, like the WD FTP site (in /pub/standards) has a copy of the ATA-2 standard as well.

Enhanced IDE Fast-ATA/ATA-2 FAQ

John Wehman and Peter den Haan

<<ftp://fission.dt.wdc.com/>>, a Western Digital FTP site, contains loads of material pertaining to the SFF Committee, the ATA, ATA-2 and ATAPI standards, the Phoenix Enhanced BIOS spec, and archives of the various reflectors (mailing lists, such as the ATA-2 and ATAPI lists). Look in /pub/standards/; the SFF stuff is in /pub/standards/sff/specs/, the x3t13 in /pub/standards/x3t13/. The mailing lists themselves can be accessed through majordomo@dt.wdc.com; send a message

To: majordomo@dt.wdc.com

help

end

to get more information.

<<ftp://ftp.symbios.com/pub/standards/>> also contains information about ATA, ATA-2, ATA-3, SCSI-2 and many more standards.

<ftp://www.ptltd.com/pub/phoenix_docs/> has many Phoenix specs such as their Enhanced BIOS document.

<<ftp://hddtech.millcomm.com>> Tech sheets (I really have to take a look there someday).

Other pointers are appreciated