# BIOS Enhanced Disk Drive Services 4.0

## T13 Technical Proposal – e08134r10

**By**
**Rob Strong – Intel Corporation**
**robert.strong@intel.com**

**2008-21~~19~~-10~~05~~**

## Document Status

| Revision History | | |
|---|---|---|
| **Rev** | **Date** | **Description** |
| 0 | June 19, 2007 | Initial draft |
| 1 | October~~Augus~~ 21~~t 19~~, 2008 | Revised based on comments from October~~April~~ Plenary Meeting |
| | | |
| | | |

# 1 Introduction

Existing BIOS Enhanced Disk Drive Services – 3, (EDD-3), specification defines an extension, “Get Device Parameters (FN 48h)”, for system firmware to return details related to device geometry, capabilities as well as physical connection for the requested device.

Although the current specification comprehensively covers the physical interfaces available at the time of the specifications writing, there have been many changes in the industry related to the manner devices are enumerated (seen) from host (legacy mode vs. native HBA specific mode) as well as to the manner devices are accessed from host (task file vs. command FIS).

With the updated information, system firmware and host operating system can make more informed decision on the device connections as well as feature set that can be provided.

# 2 Scope

This new feature is to be included in EDD revision-4

# 3 Overview

System firmware returns the information for “Get Device Parameters” extension per the proposal to enable the option ROM’s and operating systems to differentiate between SATA devices functioning in IDE-emulation mode vs. devices functioning in native-SATA mode – including those SATA devices behind a SATA port multiplier.

# 4 Changes to EDD-rev 3

## 4.1 Changes to section 8.24

### 8.24 Get Device Parameters (FN 48h)

This function returns default device parameters. It shall be mandatory regardless of the interface subset that is supported. **Table 1**~~**Table 1**~~ defines the result buffer. On entry the first word of the result buffer shall be the buffer length in bytes.

Entry:
AH = 48h
DL = BIOS device number
DS:SI = address of result buffer (See **Table 1**~~**Table 1**~~)

Exit:
carry clear
AH = 0
DS:SI = address of result buffer (See **Table 1**~~**Table 1**~~)
carry set
AH = Error code

**Table 1 - Result Buffer**

| Offset | Type | Description |
|---|---|---|
| 0 | Word | The caller shall set this value to the maximum Result Buffer length, in bytes. If the length of this buffer is less than 30 bytes, this function shall not return the pointer to Drive Parameter Table (DPT) extension. If the buffer length is 30 or greater on entry, it shall be set to 30 on exit. If the buffer length is between 26 and 29, it shall be set to 26 on exit. If the buffer length is less than 26 on entry an error shall be returned. |
| 2 | Word | Information Flags. A value of one in a bit indicates that the feature shall be available. A value of zero in a bit indicates the feature shall be not available and shall operate in a manner consistent with the conventional INT 13h interface.<br>**Bit** — **Description**<br>0 — DMA boundary errors are handled transparently<br>1 — The geometry returned in bytes 4-15 shall be valid<br>2 — Media shall be removable. Bits 4-6 are not valid if this bit is cleared to zero<br>3 — Device supports write verify<br>4 — Device has media change notification<br>5 — Media shall be lockable<br>6 — Device geometry shall be set to maximum and no media shall be present when this bit is set to one<br>7 — BIOS calls INT13h FN 50h to access the device<br>8-15 — Reserved |
| 4 | DWord | Number of default cylinders. The content of this field shall be one greater than the maximum cylinder number. INT 13h FN 08h shall be used to find the logical number of cylinders. |
| 8 | DWord | Number of default heads. The content of this field shall be one greater than the maximum head number. INT 13h FN 08h shall be used to find the logical number of heads. |
| 12 | DWord | Number of default sectors per track. The content of this field shall be the same as the maximum sector number because sector addresses are 1 based. INT 13h FN 08h shall be used to find the logical number of sectors per track. |
| 16 | QWord | Number of sectors. This shall be one greater than the maximum sector number. If this field is greater than 15,482,880 then word 2, bit 1 shall be cleared to zero. |
| 24 | Word | Number of bytes in a sector. |
| 26 | DWord | Pointer to the Device Parameter Table Extension (DPTE). This field follows the seg:offset address format. The DPTE shall only be present if INT 13h, FN 41h, CX register bit 2 is set to one. This field |

**Table 1 - Result Buffer**

| Offset | Type | Description |
|---|---|---|
| | | points to a temporary buffer that the BIOS may invalidate on subsequent INT 13h calls. If the length of this result buffer is less than 30, the DPTE shall not be present. This field is only used for INT 13h based systems configured with ATA or ATAPI devices. |
| 30 | Word | 0BEDDh - Key, indicates presence of Device Path Information |
| 32 | Byte | Length of Device Path Information including the key. The content of this byte shall be 2Ch |
| 33 | Byte | Reserved. The value in this field shall be 00h. |
| 34 | Word | Reserved. The value in this field shall be 0000h. |
| 36 | ASCII | Host bus type, 4 bytes. ASCII data shall be left justified and padded with the value 20h<br>PCI — PCI Local Bus — 50h 43h 49h 20h<br>ISA — Conventional 16 bit fixed bus — 49h 53h 41h 20h<br>PCIX — PCI-X Bus — 50h 43h 49h 58h<br>IBND — Infiniband — 49h 42h 4Eh 44h<br>XPRS — PCI Express — 58h 50h 52h 53h<br>HTPT — HyperTransport — 48h 54h 50h 54h |
| 40 | ASCII | Interface type, 8 bytes. ASCII data shall be left justified and padded with the value 20h<br>ATA — ATA/ATAPI compliant device using ATA commands — 41h 54h 41h 20h 20h 20h 20h 20h<br>ATAPI — ATA/ATAPI compliant device using ATAPI commands — 41h 54h 41h 50h 49h 20h 20h 20h<br>SCSI — SCSI compliant device — 53h 43h 53h 49h 20h 20h 20h 20h<br>USB — USB Mass Storage compliant device — 55h 53h 42h 20h 20h 20h 20h 20h<br>1394 — 1394 Mass Storage device — 31h 33h 39h 34h 20h 20h 20h 20h<br>FIBRE — Fibre Channel — 46h 49h 42h 52h 45h 20h 20h 20h<br>$I_2O$ — Intelligent Input/Output — 49h 32h 4Fh 20h 20h 20h 20h 20h<br>RAID — Redundant Array of Inexpensive Disks (RAID) member — 52h 41h 49h 44h 20h 20h 20h 20h<br>SATA — Serial ATA device in non-PATA~~IDE~~ emulation mode (Ed note: Add PATA definition as per existing standards) — 53h 41h 54h 41h 20h 20h 20h 20h<br>SAS — Serial Attached SCSI — 53h 41h 53h 20h 20h 20h 20h 20h |
| 48 | QWord | Interface Path, 8 bytes. See Table 2~~Table 2~~ for more information |
| 56 | Double QWord | Device Path. See Table 3~~Table 3~~ for more information. |
| 72 | Byte | Reserved. The value in this field shall be 00h. |
| 73 | Byte | Checksum for Device Path Information includes the 0BEDDh signature. The content of this field shall be the two's complement of the unsigned sum of offset 30 through 72. The unsigned sum of offset 30 through 73 shall be 0. |

### 8.24.1 Interface Path

The Interface Path field at offset 48 allows software external to a system BIOS to locate mass storage device interface chips. The format of this field shall be dependent on the Host Bus type, offsets 36 through 39 of the result buffer. The following formats are defined:

**Table 2 - Interface Path Definitions**

| Host Bus Type | Offset | Type | Definition |
|---|---|---|---|
| ISA | 48 | Word | 16-bit base address |
| | 50 | Word | Reserved. The value in this field shall be 0000h |
| | 52 | DWord | Reserved. The value in this field shall be 00000000h |
| PCI | 48 | Byte | PCI bus number. Values 00h through FEh shall represent a valid PCI bus. Value FFh shall indicate that this field is not used. |
| | 49 | Byte | PCI slot number. Values 00h through FEh shall represent a valid PCI slot. Value FFh shall indicate that this field is not used. |
| | 50 | Byte | PCI function number. Values 00h through FEh shall represent a valid PCI function. Value FFh shall indicate that this field is not used. |
| | 51 | Byte | Channel number. If more than one interface of the same type is accessed through a single Bus, Slot, Function, then the channel number shall identify each interface. If there is only one interface, the content of this field shall be cleared to zero. If there are two interfaces, such as an ATA Primary and Secondary interface, the primary interface shall be zero, and the secondary interface shall be one.<br>Values 00h through FEh shall represent a valid Channel Number. Value FFh shall indicate that this field is not used<br>In case of device connected to SATA controller functioning in non-PATA~~IDE~~ emulation mode, this byte shall be FFh~~0~~. |
| | 52 | DWord | Reserved. The value in this field shall be 00000000h |
| PCI-X | 48 | Byte | PCI-X bus number. Values 00h through FEh shall represent a valid PCI bus. Value FFh shall indicate that this field is not used. |
| | 49 | Byte | PCI slot number. Values 00h through FEh shall represent a valid PCI slot. Value FFh shall indicate that this field is not used. |
| | 50 | Byte | PCI function number. Values 00h through FEh shall represent a valid PCI function. Value FFh shall indicate that this field is not used. |
| | 51 | Byte | Channel number. If more than one interface of the same type is accessed through a single Bus, Slot, Function, then the channel number shall identify each interface. If there is only one interface, the content of this field shall be cleared to zero. If there are two interfaces, such as an ATA Primary and Secondary interface, the primary interface shall be zero, and the secondary interface shall be one.<br>Values 00h through FEh shall represent a valid Channel Number. Value FFh shall indicate that this field is not used<br>In case of device connected to SATA controller functioning in non-PATA~~IDE~~ emulation mode, this byte shall be FFh~~0~~. |
| | 52 | DWord | Reserved. The value in this field shall be 00000000h |

| Host Bus Type | Offset | Type | Definition |
|---|---|---|---|
| Infiniband | 48h | Qword | Reserved |
| PCI Express | 48h | Qword | See PCI |
| HyperTransport | 48h | Qword | See PCI |

### 8.24.2 Device Path

The Device Path at offset 56 combined with the Interface Path allows software external to a system BIOS to locate a specific mass storage device. The Device Path field provides a path from an interface to a specific device. The format of the Device Path at offset 40 through 47 is dependent on the Interface type. The following formats are defined in this standard:

**Table 3 - Device Path Definitions**

| Interface Type | Offset | Type | Definition |
|---|---|---|---|
| ATA | 56 | Byte | 00h = ATA Device 0, 01h = ATA Device 1 |
| | 57 | Byte | Reserved. The value in this field shall be 00h. |
| | 58 | Word | Reserved. The value in this field shall be 0000h. |
| | 60 | DWord | Reserved. The value in this field shall be 00000000h. |
| | 64 | QWord | Reserved. The value in this field shall be 0000000000000000h. |
| ATAPI | 56 | Byte | 00h = ATAPI Device 0, 01h = ATAPI Device 1 |
| | 57 | Byte | Logical Unit Number |
| | 58 | Byte | Reserved. The value in this field shall be 00h. |
| | 59 | Byte | Reserved. The value in this field shall be 00h. |
| | 60 | DWord | Reserved. The value in this field shall be 00000000h. |
| | 64 | QWord | Reserved. The value in this field shall be 0000000000000000h. |
| SCSI | 56 | Word | Physical Unit Number/SCSI ID |
| | 58 | QWord | Logical Unit Number |
| | 66 | Word | Reserved. The value in this field shall be 0000h. |
| | 68 | DWord | Reserved. The value in this field shall be 00000000h. |
| USB | 56 | QWord | 64-bit Serial Number as defined in the USB Mass Storage specifications |
| | 64 | QWord | Reserved. The value in this field shall be 0000000000000000h. |
| 1394 | 56 | QWord | 64-bit Extended Unique Identifier (EUI-64) |
| | 64 | QWord | Reserved. The value in this field shall be 0000000000000000h. |
| FIBRE | 56 | QWord | 64-bit Worldwide Identifier (WWID) |
| | 64 | QWord | Logical Unit Number |
| $I_2O$ | 56 | QWord | 64-bit Identity Tag |
| | 64 | QWord | Reserved. The value in this field shall be 0000000000000000h. |
| RAID | 56 | DWord | RAID array number of which this device is a member |
| | 60 | DWord | Reserved. The value in this field shall be 00000000h. |

**Table 3 - Device Path Definitions**

| Interface Type | Offset | Type | Definition |
|---|---|---|---|
| | 64 | QWord | Reserved. The value in this field shall be 0000000000000000h. |
| ~~SATA (IDE emulation mode)~~ | ~~56~~ | ~~Byte~~ | ~~00h = SATA Device 0, 01h = SATA Device 1~~ |
| | ~~57~~ | ~~Byte~~ | ~~Reserved. The value in this field shall be 00h.~~ |
| | ~~58~~ | ~~Word~~ | ~~Reserved. The value in this field shall be 0000h.~~ |
| | ~~60~~ | ~~DWord~~ | ~~Reserved. The value in this field shall be 00000000h.~~ |
| | ~~64~~ | ~~QWord~~ | ~~Reserved. The value in this field shall be 0000000000000000h.~~ |
| SATA ~~(AHCI mode)~~ | 56 | Byte | ~~Controller~~ Port # ~~(connector port #)~~:<br>Values 0 through 1Fh shall represent the SATA controller port number that this device is accessed through. Values 20h to FFh are reserved.<br>If device is connected directly to the SATA controller, this byte shall represent the SATA controller port number to which the device is connected to<br>If the device is behind a port multiplier then this byte shall represent the SATA controller port number to which the port multiplier is connected to. |
| | 57 | Byte | Multiplier port #:<br>Values 00h – 0Eh represent the port number values. Values 0Fh – FEh are reserved. If device is connected directly to the SATA controller, this byte shall be FFh<br>If device is behind a port multiplier, this byte shall represent the multiplier port # that this device is attached to. |
| | 58 | Word | Reserved. The value in this field shall be 0000h. |
| | 60 | DWord | Reserved. The value in this field shall be 00000000h. |
| | 64 | QWord | Reserved. The value in this field shall be 0000000000000000h. |
| SAS | 56 | QWord | 64-bit SAS address |
| | 64 | QWord | Logical Unit Number |

### 8.24.3 Device Parameter Table Extension (DPTE)

The DPTE provides hardware configuration information to applications that bypass INT 13h for addressing an ATA device.

**Table 4 - Device parameter table extension**

| Offset | Type | Description |
|---|---|---|
| 0-1 | Word | I/O port base address |
| 2-3 | Word | Control port address |
| 4 | Byte | Device register<br>bit 0-3 0<br>bit 4 ATA DEV bit<br>bit 5 1<br>bit 6 LBA enable (1 = enabled)<br>bit 7 1 |
| 5 | Byte | BIOS Vendor Specific. |
| 6 | Byte | IRQ information<br>bits 0-7 IRQ for this device |
| 7 | Byte | Block count for ATA READ/WRITE MULTIPLE commands |
| 8 | Byte | DMA information<br>bits 0-3 DMA channel<br>bits 4-7 DMA type |
| 9 | Byte | PIO information<br>bits 0-3 PIO type<br>bits 4-7 0 |
| 10-11 | Word | BIOS selected hardware specific option flags<br>bit 0 PIO accessing enabled<br>bit 1 DMA accessing enabled<br>bit 2 ATA READ/WRITE MULTIPLE accessing enabled<br>bit 3 CHS translation enabled<br>bit 4 LBA translation enabled<br>bit 5 Removable media<br>bit 6 ATAPI device<br>bit 7 32-bit transfer mode<br>bit 8 ATAPI device uses command packet interrupt<br>bits 9-10 Translation type<br>bit 11 Ultra DMA accessing enabled<br>bit 12 HPA Active<br>bit 13 48-bit LBA addressing enabled<br>bits 14-15 Reserved, shall be cleared to zero |
| 12-13 | Word | Reserved, shall be cleared to zero |
| 14 | Byte | 30h, version level of this table. |
| 15 | Byte | Checksum, 2's complement of the 8 bit unsigned sum of bytes 0-14 |

#### 8.24.3.1 Offset 0-1 - I/O port base

This word is the 16-bit address in I/O space of the data register in the ATA Command Block. Any application that provides a proprietary interface to the device may use this base address.

#### 8.24.3.2 Offset 2-3 - control port base

This word is the 16-bit address in I/O space of the device control register. Any application that provides a proprietary interface to the device may use this address.

#### 8.24.3.3 Offset 4 - head prefix

The upper four bits of this byte shall be logically ORed with the head number, or upper four bits of the LBA, each time the disk is addressed. It contains the ATA DEV bit and the LBA addressing bits that are preset, and makes these functions transparent to any software using this extension. The LBA addressing bit is set for each disk access and shall not be used to determine the LBA capability of the system. See the LBA translation enabled bit described in clause 0 for system LBA capability.

#### 8.24.3.4 Offset 5 - BIOS use only

BIOS use only.

#### 8.24.3.5 Offset 6 - IRQ number

Each ATA channel requires an assigned Interrupt number. This byte identifies which IRQ is used by this device's channel.

#### 8.24.3.6 Offset 7 - READ/WRITE MULTIPLE command block count

If the device was configured to use the READ/WRITE MULTIPLE command, then this field shall contain the block size of the transfer, in sectors, used by the BIOS.

#### 8.24.3.7 Offset 8 - DMA channel/Multiword DMA Type

If the BIOS has configured the system to perform multiword DMA data transfers in place of PIO transfers, this field shall specify the DMA mode in the upper four bits, as per the definition in ATA/ATAPI-6 or later, and the DMA Channel in the lower four bits. ATA channels that support PCI DMA bus mastering shall set the DMA channel to zero. Note that the DMA Type field does not follow the format of the data returned by the device. The value of the DMA mode shall not be limited to two.

#### 8.24.3.8 Offset 9 - PIO type

If the BIOS has configured the system to perform PIO data transfers other than mode 0, this field shall specify the PIO mode as per the definition in ATA-5 or later.

#### 8.24.3.9 Offset 10-11 - BIOS selected hardware specific option flags

These bytes specify the current hardware options enabled by the BIOS, a bit for each of the options listed below.

##### 8.24.3.9.1 Bit 0 - fast PIO

If the system is configured for a PIO mode greater than 0, this bit shall be set to one and byte 9 (PIO Type) shall be used to configure the system. If this bit is cleared to zero, the PIO-Type field shall be ignored.

##### 8.24.3.9.2 Bit 1 - fast DMA

If the system is configured for DMA, this bit shall be set to one and byte 8 (DMA Channel/DMA Type) should be used to configure the system. If this bit and bit 11, clause 0, are cleared to zero, then the DMA Channel/DMA Type field shall be ignored.

#### 8.24.3.9.3 Bit 2 - ATA READ/WRITE MULTIPLE

If the system is configured for multi-sector transfers, this bit shall be set to one and byte 7 (sector count) specifies the number of sectors used for each data transfer. If block PIO is disabled, ignore the block count field.

#### 8.24.3.9.4 Bit 3 - CHS translation

If the device reports more than 1024 cylinders in the IDENTIFY DEVICE command data, this bit shall be set to one. See clause 0 to determine the method of geometry translation.

#### 8.24.3.9.5 Bit 4 - LBA translation

If the system is configured for LBA type addressing, this bit shall be set to one and the Extended INT 13h interface (FN 41h through 48h) shall pass LBA values directly to the device. The conventional INT 13h interface shall ignore this bit and shall use CHS. LBA-type addressing shall be available on devices with less than 1024 cylinders, and therefore bit 3 (CHS translation) shall be independent from bit 4 (LBA translation).

#### 8.24.3.9.6 Bit 5 - removable media

If the device supports removable media, this bit shall be set to one and the extended INT 13h device locking and ejecting subset shall also be supported.

#### 8.24.3.9.7 Bit 6 - ATAPI device

If this ATA device implements the PACKET command feature set (ATAPI) as defined in ATA/ATAPI-5, this bit shall be set to one.

#### 8.24.3.9.8 Bit 7 - 32-bit transfer mode

If the BIOS has configured the host adapter to perform 32-bit wide data transfers, this bit shall be set to one.

#### 8.24.3.9.9 Bit 8 - ATAPI device uses command packet interrupt

If bit 6 is cleared to zero, then this field shall be ignored and shall be zero. If bit 6 is set to one, this bit indicates how the ATAPI devices signals it is ready to receive a packet command. When this bit is set to one, it indicates that the ATAPI device returns an interrupt, and sets DRQ, when it is ready for a packet. When this bit is cleared to zero, it indicates that the ATAPI device sets DRQ, without an interrupt, when it is ready for a packet.

#### 8.24.3.9.10 Bits 9-10 - translation type

If bit 3 is cleared to zero then this field shall be ignored and shall be zero. If bit 3 is set to one then this field identifies the geometric translation shown in **Table 5~~Table 5~~**.

**Table 5 - Translation Type**

| Bits 9-10 | Description |
|---|---|
| 00 | Bit-shift translation |
| 01 | LBA assisted translation |
| 10 | Reserved |
| 11 | Vendor specific translation |

#### 8.24.3.9.11Bit 11 - Ultra DMA

If the system is configured for Ultra DMA, this bit shall be set to one and byte 8 (DMA Channel/DMA Type) should be used to configure the system. If this bit and bit 1, (Bit 1 = fast DMA, clause 0) are cleared to zero, then the DMA Channel/DMA Type field shall be ignored.

#### 8.24.3.9.12Bit 12 - HPA Active

The HPA Active bit shall be set to one if the Host Protected Area (HPA) is active. The HPA is active if a volatile or non-volatile SETMAX has been issued with a value other than the NATIVE MAX address. This bit shall be cleared to zero if NATIVE MAX sectors can be accessed using read and/or write commands.

#### 8.24.3.9.13Bit 13 - 48-bit LBA addressing enabled

The 48-bit LBA addressing enabled bit shall be set to one if the INT 13h handler is using 48-bit extended commands to access the device. The 48-bit LBA addressing enabled bit shall be cleared to 0 if the INT 13h handler is not using 48-bit extended commands to access the device.

#### 8.24.3.9.14Bits 14-15 - Reserved

Shall be cleared to zero.

### 8.24.3.10 Offset 12-13 - Reserved

Shall be cleared to zero.

### 8.24.3.11 Offset 14 - table revision

The table version shall be set to 30h indicating compliance with this standard.

### 8.24.3.12 Offset 15 - checksum

This shall be the two's complement of the 8-bit unsigned sum of bytes 0 through 14. Adding bytes 0 through 15 shall in all cases produce an 8-bit result of zero.