

ATA Command Pass-Through

August 13, 2004

Revision 1

Technical Editor:

Curtis E. Stevens
Western Digital
Phone: 949-672-7933
E-Mail: Curtis.Stevens@WDC.com

Table of Contents

1	INTRODUCTION	1
2	SCOPE	1
2.1	DEFINITION OF TERMS	1
2.1.1	<i>SCSI Device</i>	1
2.1.2	<i>Bridge</i>	1
2.1.3	<i>ATA Device</i>	1
2.1.4	<i>Host</i>	1
2.2	FIELD REPRESENTATION	1
2.2.1	<i>Word</i>	2
2.2.2	<i>DWord</i>	2
3	OVERVIEW	3
4	SCSI TRANSPORTS.....	3
4.1	ISSUING AN ATA COMMAND	3
4.2	STATUS RETURN	8
4.3	ATA STATUS RETURN DESCRIPTOR	9

List of Tables

Table 1 – Short ATA Command CDB.....	4
Table 2 – ATA Command CDB	5
Table 3 – ATA Extended Command CDB	6
Table 4 – ATA Protocols	7
Table 5 – Command Successful Status Return	9
Table 6 – Command Failed Status Return	9
Table 7 – Request Sense CDB	10
Table 8 –ATA Descriptor Return	11
Table 9 – Extended ATA Descriptor Return	12

1 Introduction

Western Digital (WD) is developing devices that use SCSI for transporting commands to an ATA device. SCSI commands are documented at www.t10.org. Some of these devices are bridged from a bus that uses SCSI protocols to a device that uses PATA/SATA protocols. This proposal defines a SCSI 16 byte CDB for issuing an ATA command, and sense information to report completion status. This mechanism allows host software to tunnel through SCSI protocol bridge devices with normal ATA and Vendor specific commands using a SCSI CDB.

2 Scope

The purpose of this specification is to allow applications that are aware of ATA devices to construct SCSI CDB's that access ATA capabilities and can use SCSI pass-through mechanisms where they exist to issue the ATA commands. This proposal does not support ATAPI.

Bridge devices, drivers, or software that conform to this specification shall pass the operation requests they receive directly to the attached ATA device. The ATA fields found in the CDB shall only be passed on to the ATA device. The bridge that is passing on the ATA information shall not check the command code or any other register passed through the CDB for validity or any other purpose. If a command or any register is invalid the device shall inform the system using the error reporting techniques described in this specification.

2.1 Definition of Terms

2.1.1 SCSI Device

The part of a device that receives SCSI CDB's and returns SCSI status

2.1.2 Bridge

The part of a device that translates SCSI CDB's to ATA commands and translates ATA status to SCSI status.

2.1.3 ATA Device

The part of a device that receives and executes ATA commands and returns ATA status.

2.1.4 Host

The part of a system that generates SCSI CDB's


2.2 Field Representation

This specification documents data structures that contain Bytes, Words, and DWords. Unless otherwise specified, these fields are all represented in **Big Endian**.


2.2.1 Word

A word is a 16-bit value. This value is normally written high byte followed by low byte. Little Endian causes the word to be stored with the low byte first.

On Paper:



In Memory:



2.2.2 DWord

A DWord is a 32-bit value. This value is usually written high word first followed by low word second. Little Endian causes the low word to be stored in memory first, followed by the high word.

On Paper:



In Memory:



3 Overview

There are many transports that use SCSI Command Data Blocks (CDB's) these transports include: USB, 1394, SAS, and iSCSI. This proposal allows an ATA device to be bridged to one of these buses. If the host is aware that an ATA device is attached to a bridge then the host can issue ATA commands by using the SCSI CDB defined in this document.

Command execution is broken into three phases: Issue Command, Data Transfer, and Status. Once the CDB is constructed and issued, the bridge can use the information contained in the CDB to setup the ATA device and issue the command. A bit is provided in the CDB to differentiate between 48 bit commands that require some registers to be stored twice and 28 bit commands that require registers to be stored once. The CDB wrapper is transport dependent and contains information indicating the source or destination for the data, amount of data to be transferred etc.

Once the command has been issued to the ATA device, if the command requires data transfer then the data is transferred to or from the device following protocol requirements of the ATA device.

When data transfer is complete the ATA device returns status. This status is returned to the host. There is a special ATA request sense descriptor defined that contains all the ATA registers. If the host is asking for descriptors to be returned with the command, then all the registers are returned in the descriptor. If not, then error and status are returned as the ASC/ASCQ codes.

4 SCSI Transports

There are three phases to executing a command: Issuing the command, data transfer, and status return. Some commands do not require data transfer, so that phase can transfer 0 bytes.

4.1 *Issuing an ATA command*

When SCSI opcode 8x is received the parameters provided in offsets 1-15 are used to setup the ATA device and initiate a command. The bridge shall not check the ATA command code. The flags provided in bytes one and three, and the CDB wrapper should provide all the information necessary to execute the command. Table 2 and Table 3 show the CDB format. If the transport does not support a 16 byte CD, SCSI opcode Ax can be used to issue a 12 byte CDB. This CDB does not support extended (48 bit) commands. Table 1 shows the layout of this Short ATA Command CDB.

Table 1 – Short ATA Command CDB

Bit								
Offset	7	6	5	4	3	2	1	0
0	Operation Code = Ax							
1	Multiple.Count			Off.Line			SRST	Extend=0
2	Reserved		Dir	Protocol				
3	HWRST	DMA	D.En	LH.En	LM.En	LL.En	SC.En	F.En
4	Features (0:7)							
5	Sector Count (0:7)							
6	LBA Low (0:7)							
7	LBA Mid (0:7)							
8	LBA High (0:7)							
9	Device							
10	Command							
11	Reserved							

Table 2 – ATA Command CDB

Bit								
Offset	7	6	5	4	3	2	1	0
0	Operation Code = 8xh							
1	Multiple.Count			Off.Line			SRST	Extend=0
2	Reserved		Dir	Protocol				
3	HWRST	DMA	D.En	LH.En	LM.En	LL.En	SC.En	F.En
4	Reserved							
5	Features (0:7)							
6	Reserved							
7	Sector Count (0:7)							
8	Reserved							
9	LBA Low (0:7)							
10	Reserved							
11	LBA Mid (0:7)							
12	Reserved							
13	LBA High (0:7)							
14	Device							
15	Command							

Table 3 – ATA Extended Command CDB

Bit								
Offset	7	6	5	4	3	2	1	0
0	Operation Code = 8xh							
1	Multiple.Count			Off.Line			SRST	Extend=1
2			Dir	Protocol				
3	HWRST	DMA	D.En	LH.En	LM.En	LL.En	SC.En	F.En
4	Features (8:15)							
5	Features (0:7)							
6	Sector Count (8:15)							
7	Sector Count (0:7)							
8	LBA Low (8:15)							
9	LBA Low (0:7)							
10	LBA Mid (8:15)							
11	LBA Mid (0:7)							
12	LBA High (8:15)							
13	LBA High (0:7)							
14	Device							
15	Command							

When the Extend bit is cleared to zero a non-extended ATA command (28 bit or less) is requested. If the CDB is an ATA Command (See Table 2), the ATA parameters are loaded from offsets 5, 7, 9, 11, 13, 14 and 15. In the case of a parallel ATA device, the registers are all accessed once and the Command (data at offset 15) shall be the last register stored.

When the Extend bit is set to one an extended ATA command (48 bit) is requested (See Table 3). In the case of a parallel ATA device, bits 8:15 shall be stored in the appropriate registers first followed by bits 0:7. The Command (data at offset 15) shall be the last register stored.

If the command is a short ATA command, the Extend bit shall be cleared to zero and treated as reserved. The ATA parameters are loaded from offsets 4 through 10 of the CDB. If the Extend bit is set to one in a short ATA command the device shall report an unsupported command.

When the SRST bit is set to one, the bridge ignores all other fields with the exception of HWRST and Off.Line. The ATA device shall execute a soft reset (SRST) as defined in ATA/ATAPI-7 unless the HWRST bit is also set to one. If the HWRST bit is set to one, SRST is not performed.

Some commands can cause the ATA device to tri-state the bus. This can cause the host to see command completion before the command is actually complete. When the host

issues a command that can cause the bus to tri-state, it shall set the Off.Line field to a value that indicates the maximum number of seconds from the time a command is issued until the device status register is valid. The valid status shall be received $2 * \text{Off.Line}$ seconds after the command register is stored. This is not the time to command completion; this is the maximum time to a valid status register. The possible delays are 0 (normal case), 2, 4, 6, 8, 10, 12, and 14 seconds.

When DMA is cleared to zero indicating a PIO type transfer, Multiple.Count indicates the number of sectors transferred per interrupt. $2^{\text{Multiple.Count}}$ sectors of data are transferred before each interrupt. This field shall be set to a non-zero value only when the command is READ MULTIPLE (C4h), READ MULTIPLE EXT (29h), WRITE MULTIPLE (C5h), WRITE MULTIPLE EXT (39h), or WRITE MULTIPLE FUA EXT (CEh).

Protocol tells the bridge which protocol to use when the ATA device executes the command. Table 4 documents the protocol definitions. This field is optional and may not be implemented by some bridges. Bridges that do not implement this field may not be able to execute queued commands. Descriptions of protocols 1-13 can be found in ATA/ATAPI-7 Volume 2. The description of protocol 14 can be found in the Serial ATA II Specification: Extensions to Serial ATA 1.0a

Table 4 – ATA Protocols

Protocol	Description
0	None Specified
1	Hard Reset
2	SRST
3	Bus Idle
4	Non-data
5	PIO Data-In
6	PIO Data-Out
7	DMA
8	Packet
9	DMA Queued
10	Device Diagnostic
11	DEVICE RESET
12	UDMA Data In
13	UDMA Data Out
14	FPDMA
15-31	Reserved

The Dir bit shall be cleared to zero to indicate that this command requires data transfer from the ATA device to the host. The Dir bit shall be set to one to indicate that this command requires the host to transfer data to the ATA device. When there is no data to transfer, this bit shall be ignored by the bridge.

The host shall set F.En to one if the command requires the Features Register as a parameter.

The host shall set SC.En to one if the command requires the Sector Count Register as a parameter.

The host shall set LL.En to one if the command requires the LBA Low Register as a parameter.

The host shall set LM.En to one if the command requires the LBA Mid Register as a parameter.

The host shall set LH.En to one if the command requires the LBA High Register as a parameter.

The host shall set D.En to one if the device requires the Device Register as a parameter. The host shall set the DEV bit to 0 and the device shall ignore the DEV bit. This is required because the presence of multiple devices cannot be detected. If multiple devices are present on a single cable, the host may indicate this to the host software by reporting multiple LUNs. If there are multiple LUNs, the bridge shall set the DEV bit to the appropriate value before setting the ATA command register.

There is no enable for the command register since ATA commands are only executed after the command register is stored.

When DMA is set to one, the bridge shall transfer data to or from the ATA device using DMA. When DMA is cleared to zero, the data shall be transferred to or from the ATA device using PIO. The transfer speed uses the maximum available transfer mode of the device. The drive may not be able to transfer data when a SET FEATURES command is issued by the host to change the transfer mode (Command=EFh, Features=03h).

When the HWRST bit is set to one, the SCSI device ignores all other fields with the exception of Off.Line. The host controller shall issue a hard reset to the device. For PATA this is a Pin 1 reset, for SATA this is a COMRESET.

4.2 Status Return

When SCSI command 8x or Ax (ATA Command and Data Transfer) is issued to the SCSI device, the ATA parameters are setup using information in the CDB. When the command is executed there are two possible outcomes: the ATA ERR bit is set or the ATA ERR bit is cleared.

When the ERR bit is cleared it means the command was successfully issued. In some cases, the command may still be in process and the status indicates that the command was received but does not indicate completion. When the "50h" status occurs, the SCSI Sense Key shall be cleared to zero indicating that no error occurred. The ASC and ASCQ codes shall also be cleared to zero.

If the ERR bit is set, the command was not accepted, had invalid parameters or failed to execute successfully. The Sense Key, ASC, and ASCQ fields shall be set to indicate the type of error that occurred.

The Sense Key, ASC, and ASCQ codes shall only be persistent across Request Sense commands.

Table 5 is an example of a status return when ERR=0:

Table 5 – Command Successful Status Return

Bit								
Offset	7	6	5	4	3	2	1	0
0	Response Code = 72h or 73h							
1	Reserved				Sense Key (0h)			
2	Additional Sense Code (00h)							
3	Additional Sense Code Qualifier (00)							
4	Reserved							
5								
6								
7	Additional Length = 00h							

Table 6 is an example of a status return when a command fails.

Table 6 – Command Failed Status Return

Bit								
Offset	7	6	5	4	3	2	1	0
0	Response Code = 72h or 73h							
1	Reserved				Sense Key			
2	Additional Sense Code							
3	Additional Sense Code Qualifier							
4	Reserved							
5								
6								
7	Additional Length = 00h							

4.3 ATA Status Return Descriptor

SCSI uses the Request Sense command to return command status. Some systems use a capability called auto request. This capability returns the request sense information after each command, even if the command did not generate a check condition. Since the sense data applies to a vendor specific command, the device may return vendor specific sense data. Table 7 documents the format for the request sense command, and Table 9 defines the format of the ATA request sense data descriptor.

The SCSI device shall support the ATA descriptor. Each time the ATA descriptor is requested by the host, the SCSI device shall read the ATA registers and return the appropriate values. If the last command executed was an extended command, then 48-bit extended status is returned. If the last command executed was not extended, then 28-bit status is returned.

Table 7 – Request Sense CDB

Bit								
Offset	7	6	5	4	3	2	1	0
0	Operation Code = 03h							
1	Reserved							DESC
2	Reserved							
3	Reserved							
4	Allocation Length							
5	Control							

The ATA Request Sense Descriptor is returned when the DESC bit is set to one. If the DESC bit is cleared to zero, standard request sense information is returned as shown in Table 5 and Table 6.

Allocation Length is the size of the host buffer that shall be used to store the Request Sense response. When the DESC bit is set to one, Allocation Length should be set to a number $\geq 0Eh$. When DESC is cleared to zero, the Allocation Length should be ≥ 8 .

Table 8 –ATA Descriptor Return

		Bit							
Offset	7	6	5	4	3	2	1	0	
0	Response Code = 72h or 73h								
1	Reserved				Sense Key				
2	Additional Sense Code								
3	Additional Sense Code Qualifier								
4	Reserved								
5									
6									
7	Additional Length = 0Eh								
8	Descriptor Code = xx								
9	Additional Descriptor Length = 0Eh								
10	Reserved							Extend=0	
11	Error								
12	Sector Count (0:7)								
13	Reserved								
14	LBA Low (0:7)								
15	Reserved								
16	LBA Mid (0:7)								
17	Reserved								
18	LBA High (0:7)								
19	Reserved								
20	Device								
21	Status								

Table 9 – Extended ATA Descriptor Return

		Bit							
Offset	7	6	5	4	3	2	1	0	
0	Response Code = 72h or 73h								
1	Reserved				Sense Key				
2	Additional Sense Code								
3	Additional Sense Code Qualifier								
4	Reserved								
5									
6									
7	Additional Length = 0Eh								
8	Descriptor Code = xx								
9	Additional Descriptor Length = 0Eh								
10	Reserved							Extend=1	
11	Error								
12	Sector Count (0:7)								
13	Sector Count (8:15)								
14	LBA Low (0:7)								
15	LBA Low (8:15)								
16	LBA Mid (0:7)								
17	LBA Mid (8:15)								
18	LBA High (0:7)								
19	LBA High (8:15)								
20	Device								
21	Status								

When Extend bit is set to one, all the response data in the descriptor is valid (See Table 9). If the Extend bit is cleared to zero (See Table 8), the command did not use 48 addressing capability and the bytes labeled 8:15 shall be ignored.