

ATA Command Pass-Through

October 14, 2004

Revision 4

Technical Editor:

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

Revision History		
Date	Revision	Description
11-Aug-04	0	Initial Revision
13-Aug-04	1	<ol style="list-style-type: none"> 1. Added 5 bit protocol field. 2. Added direction. Although this is in the wrapper, some people felt that the information from the wrapper may not be available at all levels. 3. Removed statements indicating that queuing is not supported by this proposal. The addition of the protocol field allows the host to perform ATA queuing functions
16-Aug-04	2	<ol style="list-style-type: none"> 1. Added a disclaimer about SET FEATURES changing operational modes 2. Added statements to make it more clear that the Bridge owns the data transfer configuration of the ATA device. 3. Changes terms to be SCSI standard
1-Sep-04	3	<ol style="list-style-type: none"> 1. Added revision history. This was omitted from previous versions of the doc. I have also added the history for versions 0-2. 2. Turned off change-bars, they were interfering with the document restructure. 3. Added control byte to CDB's. This required a restructure of the 16-byte CDB flags. 4. Removed DMA and DIR flags as a result of #3. 5. Changed Protocol field to 4 bits as a result of #3. 6. Removed the no protocol specified field to free up a protocol number. 7. Made grammatical changes requested at the SAT WG 8. Removed the D.EN bit. All commands are defined to use the device register since it has the DEV bit for selecting a target device for the command. 9. Added a CC bit to allow the host to force a CHECK CONDITION at command completion. This is needed to prevent registers from being lost when a command completes successfully. 10. Updated overview diagram to indicate a host. 11. Deleted tables that documented request sense and sense data returns. These were causing confusion. The ATA sense data descriptors now follow documentation standards more closely. 12. Reduced the size of the Off_Line field and changed the formula to give a similar range, but with fewer

		<p>intermediate values.</p> <ol style="list-style-type: none">13. Changed more terms to match SCSI usage.14. Added a section on Inquiry that limits this proposal to peripheral device types 00h (SBC-2) and 0Eh (RBC).15. Removed the section on field representation.
2-Sep-04	4	<ol style="list-style-type: none">1. Fixed some references in the revision history for revision 3.2. Removed the reference to packet protocol from the ATA Protocol Table since this proposal does not support ATAPI.3. Modified definition of terms to only include new terms. Also added definitions of Host and Translator.4. Modified definition of Bridge.5. Added transfer length capability for both 12 and 16 bit CDB's

Table of Contents

1	INTRODUCTION.....	1
2	SCOPE.....	1
2.1	DEFINITION OF TERMS	1
2.1.1	<i>ATA Device</i>	1
2.1.2	<i>Bridge</i>	1
2.1.3	<i>Host</i>	1
2.1.4	<i>Translator</i>	2
3	OVERVIEW.....	2
4	SCSI TRANSPORTS.....	3
4.1	ISSUING AN ATA COMMAND.....	3
4.1.1	<i>BB</i>	6
4.1.2	<i>CC</i>	6
4.1.3	<i>Device</i>	7
4.1.4	<i>Extend</i>	7
4.1.5	<i>Multiple_Count</i>	7
4.1.6	<i>Off.LineOff_Line</i>	7
4.1.7	<i>Protocol</i>	8
4.1.8	<i>T_Dir</i>	8
4.1.9	<i>T_Length</i>	8
4.2	STATUS RETURN	9
4.3	ATA STATUS RETURN DESCRIPTOR	9
5	SCSI INQUIRY.....	11

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	8
Table 5 – T_Length Values.....	9
Table 6 –ATA Descriptor Return	10
Table 7 – Extended ATA Descriptor Return	10

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 translator 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 ATA device shall inform the translator using reporting techniques found in ATA/ATAPI-7. The translator shall report errors using the methods defined in this proposal.

The SCSI to ATA bridge configures the ATA host and device for the PIO, DMA, and UDMA speeds that the bridge supports. SET FEATURES commands that change the PIO/DMA modes are outside the scope of this standard.

2.1 Definition of Terms

2.1.1 ATA Device

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

2.1.2 Bridge

The part the system that receives SCSI CDB's from a SCSI Initiator, translates them to ATA commands, and then sends them to an ATA device using an ATA Host. The Bridge is also the part of the system that receives status information in response to an ATA command from an ATA Device, translates the status into sense data, and returns the sense data to a SCSI Initiator.

2.1.3 Host

The Host is the portion of the system that constructs CDB's and associated wrappers for delivery by a SCSI Initiator Port.

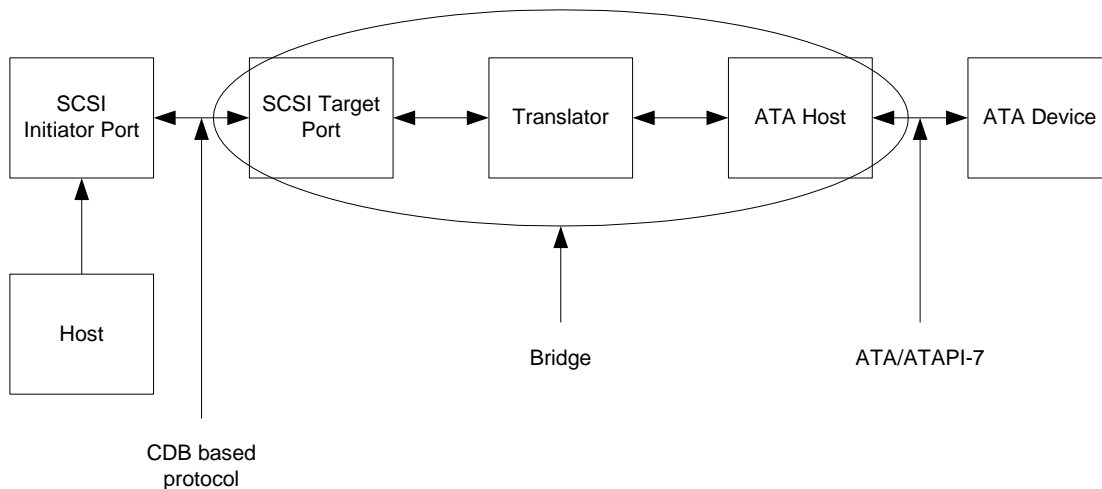
2.1.4 Translator

A Translator receives SCSI CDB's from a SCSI Target and uses these CDB's to issue commands to an ATA device using an ATA Host. A Translator also receives status from an ATA Host in response to an ATA command and uses this information to create sense data that a SCSI Target can deliver to a SCSI Initiator.

3 Overview

There are many transports that use SCSI Command Data Blocks (CDB's) these transports include: USB, 1394, SAS, and iSCSI, etc. 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.

SCSI to ATA conversion is composed of 3 parts as diagramed below:



The bridge is composed of a SCSI Target, an ATA Host, and a translator. The translator receives CDB's from the SCSI Target, converts them to ATA/ATAPI-7 commands, and issues them to the ATA device using the ATA Host. The SCSI Target may be conceptual, possibly only existing in software; or it could be any device that receives SCSI CDB's including a SAS Target, 1394 device, USB, etc. The ATA Host shall be one of the following:

1. A standard SATA ATA/ATAPI-7 compliant host
2. A standard PATA ATA/ATAPI-7 compliant host
3. A standard SAS STP Capable Transport

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. This specification allows for simple bridges that may only pass-through standard ATA commands that do not involve queuing as well as more sophisticated bridges that can transport queued commands.

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 yy is received the parameters provided in offsets 1-14 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 two of the CDB, as well as data in the CDB wrapper shall provide all the information necessary to execute the command. Table 2 and Table 3 show the CDB formats. If the transport does not support a 16 byte CDB, SCSI opcode xx 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.

The SCSI to ATA Bridge configures the ATA host and device for the PIO, DMA, and UDMA speeds that the bridge supports. SET FEATURES commands that change the PIO/DMA modes are outside the scope of this standard. If a CDB is issued that changes the PIO/DMA timings, communications may be lost with the ATA device.

Table 1 – Short ATA Command CDB

Bit								
Offset	7	6	5	4	3	2	1	0
0	Operation Code = xx							
1	Multiple_Count			Protocol				Extend=0
2	Off_Line	CC	Reserved	T_Dir	BB	T_Length		
3	Features (0:7)							
4	Sector Count (0:7)							
5	LBA Low (0:7)							
6	LBA Mid (0:7)							
7	LBA High (0:7)							
8	Device							
9	Command							
10	Reserved							
11	Control							

Table 2 – ATA Command CDB

Bit								
Offset	7	6	5	4	3	2	1	0
0	Operation Code = yy							
1	Multiple_Count			Protocol				Extend=0
2	Off_Line	CC	Reserved	T_Dir	BB	T_Length		
3	Reserved							
4	Features (0:7)							
5	Reserved							
6	Sector Count (0:7)							
7	Reserved							
8	LBA Low (0:7)							
9	Reserved							
10	LBA Mid (0:7)							
11	Reserved							
12	LBA High (0:7)							
13	Device							
14	Command							
15	Control							

Table 3 – ATA Extended Command CDB

Bit								
Offset	7	6	5	4	3	2	1	0
0	Operation Code = yy							
1	Multiple_Count			Protocol				Extend=0
2	Off_Line		CC	Reserved	T_Dir	BB	T_Length	
3	Features (8:15)							
4	Features (0:7)							
5	Sector Count (8:15)							
6	Sector Count (0:7)							
7	LBA Low (8:15)							
8	LBA Low (0:7)							
9	LBA Mid (8:15)							
10	LBA Mid (0:7)							
11	LBA High (8:15)							
12	LBA High (0:7)							
13	Device							
14	Command							
15	Control							

4.1.1 BB

The BB field shall indicate that the transfer units for T_Length field are bytes or blocks. BB=0 shall indicate that the transfer units are bytes and BB=1 shall indicate that the transfer units are blocks. This field shall be ignored when T_Length=0

4.1.2 CC

The Check Condition (CC) bit informs the bridge that the host requires register information as a part of command completion.

When the CC bit is set to one, the SCSI Target shall generate a CHECK CONDITION when the ATA command completes, even if the command completed successfully. If the command completed successfully, the SCSI Target shall set SK, ASC, and ASCQ to zero. This shall indicate that an error did not occur to the application, but it will allow the application to read the registers from the ATA Descriptor (See Table 6), or Extended ATA Descriptor (See Table 7).

When the CC bit is cleared to zero, the SCSI Target shall only generate a CHECK CONDITION when an error occurs. See section 4.2 for a description of ATA error

conditions. When the CC bit is cleared to zero, the host does not need the ATA registers when the command completes successfully.

4.1.3 Device

The host shall set the DEV bit in the device register to 0 and the device shall ignore the DEV bit. This is required because the presence of multiple devices cannot be detected using this proposal. 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 translator shall set the DEV bit to the appropriate value before issuing the ATA command.

4.1.4 Extend

When the Extend bit is cleared to zero a non-extended ATA command (28 bit or less) is requested. If the CDB is a Short ATA Command (See Table 1), the ATA parameters are loaded from offsets 4 through 10. If the CDB is an ATA Command CDB (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 3 through 9 of the CDB. If the Extend bit is set to one in a short ATA command the bridge shall report an unsupported command.

4.1.5 Multiple_Count

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).

4.1.6 Off.LineOff_Line

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.LineOff_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 $(2^{\text{Off.LineOff_Line}}) * 2 - 2$ 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, 6, and 14 seconds.

4.1.7 Protocol

Protocol tells the bridge which protocol to use when the ATA device executes the command. Table 4 documents the protocol definitions. Descriptions of protocols 0-12 can be found in ATA/ATAPI-7 Volumes 2 and 3. The description of protocol 13 can be found in the Serial ATA II Specification: Extensions to Serial ATA 1.0a.

If the protocol does not properly match the command, communications with the ATA device may be lost.

Table 4 – ATA Protocols

Protocol	Description
0	Hard Reset
1	SRST
2	Bus Idle
3	Non-data
4	PIO Data-In
5	PIO Data-Out
6	DMA
7	DMA Queued
8	Device Diagnostic
9	DEVICE RESET
10	UDMA Data In
11	UDMA Data Out
12	FPDMA
13-15	Reserved

If the host selects protocol 0, the bridge shall issue a pin 1 reset to PATA devices and COMRESET to SATA devices. When this protocol is selected, only the Protocol and Off_Line fields are valid, the translator shall ignore all other fields in the CDB.

If the host selects protocol 1, the bridge shall issue a soft reset as defined in ATA/ATAPI-7. When this protocol is selected, only the Protocol and Off_Line fields are valid, the translator shall ignore all other fields in the CDB.

4.1.8 T_Dir

T_Dir shall indicate the direction of data transfer. T_Dir=0 shall indicate that data is transferred from the initiator to the target and T_Dir=1 shall indicate that data is transferred from the target to the initiator. This field shall be ignored when T_Length=0.

4.1.9 T_Length

T_Length specifies source for the transfer length of the command as described in Table 5. The transfer length shall be an 8 or 16 bit number. The value of the Extend bit (see 4.1.4) shall determine the number of bit in the transfer length. See 4.1.1 for a description on how to determine if the unit of transfer is a byte or a block.

Table 5 – T_Length Values

Value	Description
00h	No data is transferred
01h	Transfer Length is found in the Feature field.
02h	Transfer Length is found in the Sector Count field
03h	Transfer Length is found in the command wrapper

4.2 Status Return

When SCSI command xx or yy (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 or DF bit are set or both bits are clear.

When the ERR and DF bits are 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 both the ERR and DF bits are cleared, 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 or DF 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.

Some ATA commands return information in the registers. Requesting the ATA Descriptor using Request Sense shall be used retrieve this information.

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

4.3 ATA Status Return Descriptor

SCSI uses the Request Sense command to return command status. Table 6 and Table 7 define the format of the ATA request sense data descriptor.

The SCSI device shall support the ATA descriptor if it accepts the ATA Command and/or ATA Extended Command CDB's documented earlier. Each time the ATA descriptor is requested by the host, the translator 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 6 –ATA Descriptor Return

Bit								
Offset	7	6	5	4	3	2	1	0
0	Descriptor Code = zz							
1	Additional Descriptor Length = 0Eh							
2	Reserved							Extend=0
3	Error							
4	Reserved							
5	Sector Count (0:7)							
6	Reserved							
7	LBA Low (0:7)							
8	Reserved							
9	LBA Mid (0:7)							
10	Reserved							
11	LBA High (0:7)							
12	Device							
13	Status							

Table 7 – Extended ATA Descriptor Return

Bit								
Offset	7	6	5	4	3	2	1	0
0	Descriptor Code = zz							
1	Additional Descriptor Length = 0Eh							
2	Reserved							Extend=1
3	Error							
4	Sector Count (8:15)							
5	Sector Count (0:7)							
6	LBA Low (8:15)							
7	LBA Low (0:7)							
8	LBA Mid (8:15)							
9	LBA Mid (0:7)							
10	LBA High (8:15)							
11	LBA High (0:7)							
12	Device							
13	Status							

When Extend bit is set to one, all the response data in the descriptor shall be valid (See Table 7). If the Extend bit is cleared to zero (See Table 6), the previous command did not use 48 addressing capability.

5 SCSI Inquiry

The SCSI Inquiry command returns information describing the device type. Bridges that conform to this proposal return device type return Standard Inquiry Data consistent with an ATA mass storage device. Bridges conforming to this proposal shall return a peripheral device type of 00h or 0Eh. There is no indication to the Host that this is an ATA device attached to a SCSI Target. The Host may determine support by issuing an ATA (Extended) Command CDB and checking for success.