

ACS-2 Synchronize with SATA 2.6

February 11, 2009

Revision 2

Technical Editor:

Jim Hatfield
389 Disc Drive
Longmont, CO 80503
720-684-2120
James.C.Hatfield@Seagate.com

Document Status

Revision History		
Rev	Date	Description
0	November 1, 2008	1) Initial Revision
1	February 2, 2009	1) Updated the changes for IDENTIFY DEVICE and the NCQ Command Error log 2) Removed material that is not to be changed.
2	February 11, 2009	1) Add request for assignment of a Major Transport Revision id for SATA Rev. 3.0

1 Introduction

ATA8-ACS incorporated some material from the SATA 2.5 specification. The SATA-IO specification has added new features and functions resulting in the SATA 2.6 specification.

2 .Scope

T13 has a copyright release from SATA-IO to include selected sections into ACS-2.

e08016r0-SATA-IO-INCITS_ACS-2_Copyright_License_rev_0_with_Exhibits.DOC

In addition, the following SATA-IO Errata against the Serial ATA 2.6 specification are to be included (for the permitted sections):

ECN # 004 - ATA Log & Subcode reservations

ECN # 013 - Section 13 Corrections

ECN # 022 - Editorial cleanup - 8KB, IDENTIFY DEVICE

3 Overview

SATA 2.6 added a Priority input to the READ FPDMA QUEUED and WRITE FPDMA QUEUED commands, and defined updated requirements for IDLE IMMEDIATE with the unload feature.

This proposal also assigns a bit to indicate SATA Rev. 3.0 in the Major Transport Revision.

This proposal updates these sections of ACS-2:

- a) 4.15 Native Command Queuing (NCQ) feature set
- b) 7.16 IDENTIFY DEVICE
- c) 7.28 READ FPDMA QUEUED
- d) 7.48 SET FEATURES
- e) 7.68 WRITE FPDMA QUEUED
- f) A.10 NCQ Command Error Log

4 Changes to ACS-2

4.1 Native Command Queuing (NCQ) feature set (4.15)

4.1.1 Overview

The optional NCQ feature set provides support for devices that implement the Serial Transport (see ATA8-AST). The NCQ feature set allows commands within this feature set to be accepted even though the device has not reported command completion for one or more previously accepted commands in the NCQ feature set. A device reports command completion for commands in the NCQ feature set by returning a transport dependent indicator (see ATA8-AST). The following commands are mandatory for devices implementing the NCQ feature set:

- a) READ FPDMA QUEUED
- b) WRITE FPDMA QUEUED

Devices that report support for the NCQ feature set shall also report support for the GPL feature set (see 4.10), the General Purpose Log Directory log and the NCQ Command Error log.

If the device receives a command that is not an NCQ command while NCQ commands are in the queue, then the device shall return command aborted for the new command and for all of the NCQ commands that are in the queue.

All the commands in the NCQ feature set shall include a NCQ Tag. If the value of the NCQ Tag exceeds the value returned in IDENTIFY DEVICE data word 75 (see 7.16.7.31), then the device shall return command aborted for the new command and for all NCQ commands that are in the queue. If the device receives an NCQ

command with a NCQ Tag value that is identical to the NCQ Tag value for another NCQ command in the queue, then the device shall return command aborted for the new command and for all the NCQ commands that are in the queue.

NOTE 1 — The NCQ Tag identifies return information (i.e., error status, data transfer and command completion).

If an error occurs while the device is processing an NCQ command, then the device shall return command aborted for all NCQ commands that are in the queue and shall return command aborted for any new commands, except a READ LOG EXT command requesting log address 10h, until the device completes a READ LOG EXT command requesting log address 10h (i.e., reading the NCQ Command Error log) without error.

NOTE 2 — The NCQ feature set uses 48-bit addresses and is intentionally not included in the 48-bit address feature set.

4.1.1.1 Priority

Host knowledge of I/O priority may be transmitted to the device as part of the command. There are two priority values for NCQ commands, normal and high. When the host marks an NCQ command as high priority, the host is requesting a better quality of service for that command than commands issued with normal priority.

The classes are forms of soft priority. The device may choose to complete a normal priority command before an outstanding high priority command, although preference should be given to the high priority commands. One example where a normal priority command may be completed before a high priority command is when the normal priority command is a cache hit, whereas the high priority command requires access of the device media.

The priority class is specified in the PRIO bit for NCQ commands (READ FPDMA QUEUED and WRITE FPDMA QUEUED). This bit may specify either the normal priority or high priority value. If a command is marked by the host as high priority, the device should attempt to provide better quality of service for the command. It is not required that devices process all high priority requests before satisfying normal priority requests.

4.1.1.2 Unload

When using Native Command Queuing ~~in a laptop environment~~, the host ~~needs~~ may need to be able to park the heads. ~~due to excessive movement (e.g. the laptop being dropped). This section defines a mechanism that the host may use to park the heads when NCQ commands are outstanding in the device. The typical time for completion of the unload operation is defined in ATA/ATAPI-7 clause 6.20.10.~~

When NCQ commands are outstanding, the device ~~is able to~~ may accept the IDLE IMMEDIATE command with the Unload ~~feature~~ Feature ~~defined in ATA/ATAPI-7 clause 6.20.~~ Upon ~~reception~~ acceptance of this command with the Unload ~~feature~~ Feature specified, the device shall:

- 1) ~~u~~Unload/park the heads immediately. ;
- 2) ~~Respond to the host with a Register—Device to Host FIS with the ERR bit set to one in the Status register~~ return command aborted since this is a non-queued command.

When the host receives the error indication, it should proceed to do a READ LOG EXT command for the NCQ Command Error log ~~log page 10h~~. In the log page, the device shall indicate whether the error was due to ~~receiving~~ accepting an UNLOAD Unload and whether the UNLOAD Unload was executed. The device shall not load the heads to the media when satisfying the READ LOG EXT command for the NCQ Command Error log ~~page 10h~~.

The READ LOG EXT command for the NCQ Command Error log ~~log page 10h~~ indicates whether the device has accepted the Unload and is in the process of executing the command. To get a definitive indication of Unload

completion (and success), the IDLE IMMEDIATE command with the Unload feature needs to be issued again after the READ LOG EXT command for [the NCQ Command Error log log page 10h](#) is executed. After the READ LOG EXT command for [the NCQ Command Error log log page 10h](#) is executed, there are no NCQ commands outstanding and the NCQ error is cleared such that the IDLE IMMEDIATE command with the Unload Feature will be processed normally and a successful status will be returned when the unload process completes successfully.

There may be a delay in issuing the IDLE IMMEDIATE command with the Unload feature to the device if the device is currently performing a data transfer for a previously issued NCQ command. If the device happens to be executing extensive data error recovery procedures, this delay could be longer than acceptable. However, this same issue may occur when a non-queued data command is outstanding and the device is performing error recovery procedures.

4.1.1.3 Intermixing Non-Native Queued Commands and Native Queued Commands

The host shall not issue a non-native queued command while a native queued command is outstanding. Upon receiving a non-native queued command while a native queued command is outstanding, the device shall signal the error condition to the host by transmitting a Register FIS to the host with the ERR and ABRT bits set to one and the BSY bit cleared to zero in the Status field of the FIS and halt command processing as described in section 13.5.1.4 except as noted below. Non-native queued commands include all commands other than the READ FPDMA QUEUED and WRITE FPDMA QUEUED commands defined in section 13.5.3.1 and section 13.5.3.2. Reception of the READ LOG EXT command with a specified log page of 10h after an error has occurred shall cause any outstanding Serial ATA native queued commands to be aborted, and the device shall perform necessary state cleanup to return to a state with no commands pending. The device shall clear all bits in the SActive register by transmitting a Set Device Bits FIS to the host with all the bits in the SActive field set to one (i.e. FFFFFFFFh). After completing a READ LOG EXT command with a specified log page of 10h, the device shall be prepared to execute subsequently issued queued commands regardless of any previous errors on a queued command. In the case that a READ LOG EXT command with a log page of 10h is issued while a native queued command is outstanding AND no error was previously reported by the device, the device shall signal an error condition. The receipt of this command when no error is outstanding shall be handled as any other non-native queued command when a native queued command is outstanding. In this case, a subsequent READ LOG EXT command with log page of 10h is required to recover from the error.

4.1.2 Command Phases

4.1.2.1 Command Acceptance

The device receives a command in the NCQ feature set and returns command acceptance. Once the device reports command acceptance, it may then accept additional commands in the NCQ feature set.

4.1.2.2 Data transmission

Data transfer should occur after acceptance of the command.

4.1.2.3 Command completion

When the transfer of all of the data requested by one or more NCQ commands occurred without error, the device returns a transport dependent indicator (see ATA8-AST) that informs the host of completion for one or more NCQ commands.

If an error occurs while processing an NCQ command, then the device shall return command aborted for the command in error and for all other NCQ commands that are in the queue. The condition of the data for any NCQ command for which a device reports command aborted is indeterminate.

4.2 IDENTIFY DEVICE - ECh, PIO Data-In

4.2.1 Feature Set

This 28-bit command is mandatory for all devices.

4.2.2 Description

The IDENTIFY DEVICE command specifies that the device shall send a 512-byte block of data to the host. See 4.2.7 for a description of the return data.

Some devices may have to read the media in order for all applicable IDENTIFY DEVICE data fields to be valid.

The IDENTIFY DEVICE data contains information regarding optional feature or command support. If the host issues a command that is indicated as not supported in the IDENTIFY DEVICE data, the device shall return command aborted for the command.

4.2.3 Inputs

Name	Description
Feature	N/A
Count	N/A
LBA	N/A
Device	<p>Bit Description</p> <p>7 Obsolete</p> <p>6 N/A</p> <p>5 Obsolete</p> <p>4 Transport Dependent - See 6.2.11</p> <p>3:0 Reserved</p>
Command	7:0 ECh

4.2.4 Normal Outputs for ATA devices

See table 99.

4.2.5 Normal Outputs for ATAPI devices

In response to this command, ATAPI devices shall return command aborted and place the PACKET feature set signature in the appropriate fields (see table 104).

4.2.6 Error Outputs

ATA devices shall not report an error.

4.2.7 Input From the Device to the Host Data Structure

4.2.7.1 Overview

Table 1 specifies the format of the IDENTIFY DEVICE data.

Table 1 — IDENTIFY DEVICE data

Word	O M	S P	F V	Description																		
76	O	S		Serial ATA Capabilities 15:13 Reserved for Serial ATA 12 1 = Supports NCQ priority information 11 1 = Supports Unload while NCQ commands are outstanding F 10 1 = Supports Phy Event Counters F 9 1 = Supports receipt of host initiated power management requests F 8 1 = Supports the NCQ feature set 7:3 Reserved for Serial ATA F 2 1 = Supports SATA Gen2 Signaling Speed (3.0Gb/s) F 1 1 = Supports SATA Gen1 Signaling Speed (1.5Gb/s)																		
222	M	B		Transport major version number (see 7.18.7.86) 0000h or FFFFh = device does not report version F 15:12 Transport Type 0h = Parallel 1h = Serial 2h-Fh = Reserved <table style="width: 100%; border: none;"> <tr> <td style="text-align: center;">Parallel</td> <td style="text-align: center;">Serial</td> </tr> <tr> <td>11:6 Reserved</td> <td>Reserved</td> </tr> <tr> <td>5 Reserved</td> <td>SATA Rev 3.0</td> </tr> <tr> <td>41:5 Reserved</td> <td>Reserved</td> </tr> <tr> <td>F 4 Reserved</td> <td>SATA Rev 2.6</td> </tr> <tr> <td>F 3 Reserved</td> <td>SATA Rev 2.5</td> </tr> <tr> <td>F 2 Reserved</td> <td>SATA II: Extensions</td> </tr> <tr> <td>F 1 ATA/ATAPI-7</td> <td>SATA 1.0a</td> </tr> <tr> <td>F 0 ATA8-APT</td> <td>ATA8-AST</td> </tr> </table>	Parallel	Serial	11:6 Reserved	Reserved	5 Reserved	SATA Rev 3.0	41:5 Reserved	Reserved	F 4 Reserved	SATA Rev 2.6	F 3 Reserved	SATA Rev 2.5	F 2 Reserved	SATA II: Extensions	F 1 ATA/ATAPI-7	SATA 1.0a	F 0 ATA8-APT	ATA8-AST
Parallel	Serial																					
11:6 Reserved	Reserved																					
5 Reserved	SATA Rev 3.0																					
41:5 Reserved	Reserved																					
F 4 Reserved	SATA Rev 2.6																					
F 3 Reserved	SATA Rev 2.5																					
F 2 Reserved	SATA II: Extensions																					
F 1 ATA/ATAPI-7	SATA 1.0a																					
F 0 ATA8-APT	ATA8-AST																					
Key:				O/M – Mandatory/optional requirement.																		
F/V – Fixed/variable content				M – Support of the word is mandatory.																		
F – The content of the field is fixed and does not change. The DCO command may change the value of a fixed field.				O – Support of the word is optional.																		
V – The contents of the field is variable and may change depending on the state of the device or the commands processed by the device.				S/P – Content applies to Serial or Parallel transport																		
X – The fixed or variable type of this field is not defined in this standard.				S – Serial Transport																		
				P – Parallel Transport																		
				B – Both Serial and Parallel Transports																		
				N – Belongs to a transport other than Serial or Parallel																		

4.2.7.2 Word 76: Serial ATA Capabilities

Word 76 indicates the capabilities of a SATA device. A PATA device shall set word 76 to 0000h or FFFFh. If word 76 is set to 0000h or FFFFh, then the device does not claim compliance with the Serial ATA specification and words 76 through 79 are not valid and shall be ignored.

If word 76 is not set to 0000h or FFFFh, then the device claims compliance with the Serial ATA specification, and words 77 through 79 shall be valid.

Bits (15:~~4~~13) of word 76 are reserved for Serial ATA.

If bit 12 of word 76 is set to one, then the device supports the Priority field in the READ FPDMA QUEUED and WRITE FPDMA QUEUED commands and optimization based on this information. This bit shall only be set to one if the device supports NCQ as shown in bit 8 of Word 76.

If bit 11 of word 76 is set to one, then the device supports performing an unload/park of the heads upon reception of the IDLE IMMEDIATE command with the Unload Feature specified while NCQ commands are outstanding. This bit shall only be set to one if the device supports NCQ as shown in bit 8 of Word 76.

If bit 10 of word 76 is set to one, then the device supports the SATA Phy Event Counters log (see A.12).

If bit 9 of word 76 is set to one, then the device supports Partial and Slumber interface power management states when initiated by the host (see SATA 2.6).

If bit 8 of word 76 is set to one, then the device supports the NCQ feature set.

Bits (7:3) of word 76 are reserved for Serial ATA.

If bit 2 of word 76 is set to one, then the device supports the Gen2 signaling rate of 3.0 Gb/s (see SATA 2.6).

If bit 1 of word 76 is set to one, then the device supports the Gen1 signaling rate of 1.5 Gb/s (see SATA 2.6).

Bit 0 of word 76 shall be cleared to zero.

4.3 READ FPDMA QUEUED - 60h, DMA Queued (7.28)

4.3.1 Feature Set

This 48-bit command is mandatory for devices implementing the NCQ feature set.

4.3.2 Description

This command requests data to be transferred from the device to the host.

4.3.3 Inputs

4.3.3.1 Overview

Name	Description
Feature	The number of logical sectors to be transferred. A value of 0000h indicates that 65,536 logical sectors are to be transferred
Count	<p>Bit Description</p> <p>15:8 PRIO 0 = Normal priority 1 = High priority</p> <p>14:8 Reserved 7:3 NCQ Tag - See 6.5.2 2:0 N/A</p>
LBA	LBA of first logical sector to be transferred
Device	<p>Bit Description</p> <p>7 FUA - See 4.3.3.3 6 Shall be set to one 5 Reserved 4 Shall be set to zero 3:0 Reserved</p>
Command	7:0 60h

4.3.3.2 Priority (PRIO)

[The Priority \(PRIO\) value shall be assigned by the host based on the priority of the command issued. The device shall make a best effort to complete High priority requests in a more timely fashion than Normal priority requests.](#)

4.3.3.3 Forced Unit Access (FUA)

When the FUA bit is set to one the device shall retrieve the data from the non-volatile media regardless of whether the device holds the requested information in its volatile cache. If the device holds a modified copy of the requested data as a result of having volatile cached writes, the modified data shall be written to the non-volatile media before being retrieved from the non-volatile media as part of this operation. When the FUA bit is cleared to zero the data shall be retrieved either from the device's non-volatile media or cache.

4.3.4 Command Acceptance Outputs

See table 116

4.3.5 Normal Outputs

See table 117.

4.3.6 Error Outputs

This return indicates that the command was aborted due to LBA out of range, a duplicate tag number, an invalid tag number, or an Interface CRC error, see table 157 for more information.

Errors that occur during the processing of this command are reported by returning a transport dependent indicator with additional information available in the NCQ Command Error log. The validity of the data transferred is indeterminate. See table 159.

4.4 SET FEATURES - EFh, Non-Data

4.4.1 Enable/Disable SATA feature

4.4.1.1 Overview

Subcommand codes 10h and 90h allow the host to enable or disable Serial ATA features. The Count field contains the specific Serial ATA feature to enable or disable. The specific Serial ATA features [in which SET FEATURES is applicable](#) are defined ~~as-defined~~ in table 2.

Table 2 — SATA Features

Count	Description
00h	Reserved for Serial ATA
01h	Non-zero Buffer Offsets
02h	DMA Setup FIS Auto-Activate optimization
03h	Device-initiated interface power state transitions
04h	Guaranteed In-Order Data Delivery
05h	Asynchronous Notification
06h	Software Settings Preservation
07h-FFh	Reserved for Serial ATA

4.5 WRITE FPDMA QUEUED - 61h, DMA Queued (7.68)

4.5.1 Feature Set

This 48-bit command is mandatory for devices implementing the NCQ feature set.

4.5.2 Description

This command causes data to be transferred from the host to the device.

4.5.3 Inputs

4.5.3.1 Overview

Name	Description
Feature	The number of logical sectors to be transferred. A value of 0000h indicates that 65,536 logical sectors are to be transferred
Count	<p>Bit Description</p> <p>15 PRIO</p> <p>14:8 Reserved</p> <p>7:3 NCQ Tag - See 6.5.2</p> <p>2:0 Reserved</p>
LBA	LBA of first logical sector to be transferred
Device	<p>Bit Description</p> <p>7 FUA - See 4.5.3.4</p> <p>6 Shall be set to one</p> <p>5 Reserved</p> <p>4 Shall be set to zero</p> <p>3:0 Reserved</p>
Command	7:0 61h

[4.5.3.2 Priority \(PRIO\)](#)

[4.5.3.3 The Priority \(PRIO\) value shall be assigned by the host based on the priority of the command issued. The device shall make a best effort to complete High priority requests in a more timely fashion than Normal priority requests.](#)

4.5.3.4 Forced Unit Access (FUA)

When the FUA bit is set to one regardless of whether volatile and/or non-volatile write caching in the device is enabled or not, the user data shall be written to non-volatile media before command completion is reported. When the FUA bit is cleared to zero the device may return command completion before the data is written to the non-volatile media.

4.5.4 Command Acceptance Outputs

See table 116.

4.5.5 Normal Outputs

See table 117.

4.5.6 Error Outputs

If the Error bit is set to one, then the device aborted the command due to LBA out of range, a duplicate tag number, an invalid tag number, or an Interface CRC error, see table 157 for more information.

Errors that occur during the processing of this command are reported by returning a transport dependent indicator with additional information available in the NCQ Command Error log, see table 158 for more information.

0.1 NCQ Command Error (Log Address 10h) (A.10)

0.1.1 Overview

The NCQ Command Error log describes the most recent NCQ command failure, is one page in length, and is defined in table 0.1. Devices supporting the NCQ feature set shall support log address 10h (i.e., NCQ Command Error).

Table 0.1 — NCQ Command Error Log

Byte	Description
0	Bit Name 7 NQ (see 0.1.3) 6 UNL 6:5 Reserved 4:0 NCQ Tag (see 0.1.2)
1	Reserved
2	Status
3	Error
4	LBA (7:0)
5	LBA (15:8)
6	LBA (23:16)
7	Device
8	LBA (31:24)
9	LBA (39:32)
10	LBA (47:40)
11	Reserved
12	Count (7:0)
13	Count (15:8)
14-255	Reserved
256-510	Vendor Specific
511	Checksum

0.1.2 NCQ Tag

If the NQ bit is cleared, the NCQ Tag field contains the NCQ Tag corresponding to the NCQ command that failed.

0.1.3 NQ

If set indicates that the error condition was a result of a non-NCQ command having been issued and that the NCQ Tag field is therefore not valid. If cleared indicates that the NCQ Tag field is valid and that the error condition applies to a NCQ command.

[0.1.4 UNL](#)

[If set in to one indicates that the error condition was a result of receiving an IDLE IMMEDIATE command with the Unload Feature specified. If cleared to zero, the reason for the error was not due to reception of an IDLE IMMEDIATE command with the Unload Feature specified. If the last command received was an Unload Immediate, the device shall not load the heads to the media when executing the READ LOG EXT command for log page 10h.](#)

If set to one, the NQ bit shall also be set to one to indicate the failure was due to reception of a non-queued command. When set to one, the value of the Status, Error, and LBA Low fields (bytes 3-5) in the log page shall be set as follows:

Status: BSY bit shall be cleared to zero and ERR bit shall be set to one

Error: ABRT bit shall be set to one

LBA [7:0]: Shall be set to C4h if the unload is being executed or has completed successfully.

Shall be set to 4Ch if the unload was not accepted or has failed.

0.1.5 Return Fields

The Status, Error, LBA and Count fields indicate the error that caused the device to stop processing NCQ commands.

NOTE 3 — The value returned in the Error field of the NCQ Command Error log may be different than the value returned in the Error field of the command Error Output structure when the initial error condition is signaled. The Error field in command Error Output structure is used for the purpose of signaling an error for an NCQ command, while the value in the Error field of the NCQ Command Error log provides specific information about the error condition.

0.1.6 Checksum

The data structure checksum is the two's complement of the sum of the first 511 bytes in the data structure. Each byte shall be added with 8-bit unsigned arithmetic and overflow shall be ignored. The sum of all 512 bytes of the data structure shall be zero.