

**Proposed  
Draft**

**Serial ATA  
International Organization**

**Version 02  
May 23, 2011**

---

**SATA31\_TPR\_D141\_20110523\_V02**

**Title: NCQ Autosense**

**Proposed change, new functionality, or behavior to Serial ATA Revision 3.1**

This is an internal working document of the Serial ATA International Organization. As such, this is not a completed standard and has not been approved. The Serial ATA International Organization may modify the contents at any time. This document is made available for review and comment only.

Permission is granted to the Promoters, Contributors and Adopters of the Serial ATA International Organization to reproduce this document for the purposes of evolving the technical content for internal use only without further permission provided this notice is included. All other rights are reserved and may be covered by one or more Non Disclosure Agreements including the Serial ATA International Organization participant agreements. Any commercial or for-profit replication or republication is prohibited. Copyright © 2005-2009 Serial ATA International Organization. All rights reserved.

This Draft Specification is NOT the final version of the Specification and is subject to change without notice. A modified, final version of this Specification ("Final Specification") when approved by the Promoters will be made available for download at this Web Site: <http://www.serialata.org>.

THIS DRAFT SPECIFICATION IS PROVIDED "AS IS" WITH NO WARRANTIES WHATSOEVER, INCLUDING ANY WARRANTY OF MERCHANTABILITY, NON-INFRINGEMENT, FITNESS FOR ANY PARTICULAR PURPOSE OR ANY WARRANTY OTHERWISE ARISING OUT OF ANY PROPOSAL, SPECIFICATION, OR SAMPLE. Except for the right to download for internal review, no license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted or intended hereunder.

THE PROMOTERS DISCLAIM ALL LIABILITY, INCLUDING LIABILITY FOR INFRINGEMENT OF ANY PROPRIETARY RIGHTS, RELATING TO USE OF INFORMATION IN THIS DRAFT SPECIFICATION. THE PROMOTERS DO NOT WARRANT OR REPRESENT THAT SUCH USE WILL NOT INFRINGE SUCH RIGHTS.

THIS DOCUMENT IS AN INTERMEDIATE DRAFT FOR COMMENT ONLY AND IS SUBJECT TO CHANGE WITHOUT NOTICE.

\* Other brands and names are the property of their respective owners.

Copyright © 2005-2011 Serial ATA International Organization. All rights reserved.

## Author Information

Author Name	Company	e-mail address
Jim Hatfield	Seagate Technology	james.c.hatfield@seagate.com
Mark Overby	nVidia	moverby@nvidia.com

## Workgroup Chair Information

Workgroup	Chairperson Name	e-mail address
Digital	Jim Hatfield	<a href="mailto:james.c.hatfield@seagate.com">james.c.hatfield@seagate.com</a>

## Document History

Version	Date	Comments
00	02/25/2011	Initial draft
<a href="#">01</a>	<a href="#">03/17/2011</a>	<ul style="list-style-type: none"> <li>a) <a href="#">Updated the title</a></li> <li>b) <a href="#">added IDENTIFY DEVICE</a></li> <li>c) <a href="#">added DCO SET</a></li> <li>d) <a href="#">added Software Settings Preservation</a></li> </ul>
<a href="#">02</a>	<a href="#">05/23/2011</a>	<ul style="list-style-type: none"> <li>a) <a href="#">Resolved TBDs</a></li> <li>b) <a href="#">added SPC-4 to 'Other references'</a></li> </ul>

## 1 Introduction

This proposal integrates NCQ commands with the ACS-2 Sense Data Reporting feature set.

## 2 Summary of the problem

The ACS-2 standard introduced the Sense Data Reporting feature set which defines a mechanism for an ATA device to return a SCSI-like sense key and sense code.

If the feature is enabled and the device needs to report an error, the device sets the Sense Data Available bit in the Status field when returning command completion with error. The host should see the bit set and respond by sending a REQUEST SENSE DATA EXT command to the device as the NEXT command. If the next command is not REQUEST SENSE DATA EXT, then the device clears the error information.

After an NCQ error, the next command processed must be one to read log 10h, or the command is ignored and the device keeps waiting for the log read request.

There is a conflict here: both require a different command to be the next command after an error is reported.

ACS-2 attempted to put in an exception for NCQ errors, but the text (bullet b) has double negatives and is confusing.

The sense data shall be cleared after:

- a) receiving any reset;
- b) acceptance of a command other than REQUEST SENSE DATA EXT command that does not read the NCQ Command Error log; or
- c) completion of a REQUEST SENSE DATA EXT command.

This proposal puts the sense data into log 10h, so that the host does not have to issue a separate REQUEST SENSE DATA EXT command after an NCQ error. Reading log 10h is embedded in some hardware implementations, so it is not practical to change the requirements for reading it as the next command.

### 3 Proposed changes

[editor note: Existing text is black. New text is marked as underlined in blue color. Material to be deleted ~~is red with strikethrough markings.~~ ]

Changes are proposed for:

- a) Other References
- b) IDENTIFY DEVICE
- c) Device Configuration Overlay Set
- d) Software Settings Preservation
- e) Queued Error Log

[editors note: Add SPC-4 to section 3.2 ]

#### 3.2 References under development

The following ANSI standards under development are referenced. Draft versions of these standards are available from <http://www.T10.org> or <http://www.T13.org>.

ATA/ATAPI-8 Serial Transport (ATA8-AST) [ANSI INCITS T13/1697-D]  
ATA/ATAPI-8 Parallel Transport (ATA8-APT) [ANSI INCITS T13/1698-D]  
ATA Host Adapter Standards - 2 (HBA-2) [ANSI INCITS T13/2014D]  
[SCSI Primary Commands - 4 \(SPC-4\) \[ANSI INCITS T10/1731-D\]](#)

[editors note: Add a new bit to IDENTIFY DEVICE (section 13.2) as noted below. Consider using word 77 bit 6 (which is the next 'reserved' bit) ]

#### 13.2 IDENTIFY DEVICE

13.2.1.18 Word 78: Serial ATA features supported

.....

[word 78 bit 7](#)

[Bit 7, when set to one indicates that the device supports NCQ Autosense \(see 13.7.3\). This bit shall only be set to one if the device supports NCQ as shown in bit 8 of Word 76.](#)

.....

Bits ~~87~~-15 are reserved and shall be cleared to zero

[editors note: Add this new text to DCO SET, word 8]

#### 13.4.2 Device Configuration Overlay Set

.....

WORD 8: Serial ATA command / feature sets supported

This word enables configuration of command sets and feature sets.

If bit 0 of word 8 is cleared to zero, then the device shall:

- a) disable support for Native Command Queuing;
- b) clear word 76 bits 8,11,and 12 in the IDENTIFY DEVICE data to zero;
- c) clear word 78 bit 7 in the IDENTIFY DEVICE data to zero;
- d) clear word 78 bits 1, 2, and 4 in the IDENTIFY DEVICE data to zero;
- e) clear word 79 bits 1, 2 and 4 in the IDENTIFY DEVICE data to zero; and
- f) if NCQ is disabled and READ FPDMA QUEUED or WRITE FPDMA QUEUED is issued to the device, the device shall abort the command with the ERR bit set to one in the Status field and the ABRT bit set to one in the Error field.

The setting of this bit is applicable to non-PACKET devices only.

[editors note: Alphabetize the list of SET FEATURES subcommands ]

[editors note: Add [new text](#) to Software Settings Preservation]

### 13.5 Software Settings Preservation (Optional)

When a device is enumerated, software configures the device using SET FEATURES and other commands. These software settings are often preserved across software reset but not necessarily across COMRESET. In Parallel ATA, only commanded hardware resets may occur, thus legacy mode software only reprograms settings that are cleared for the particular type of reset it has issued. In Serial ATA, COMRESET is equivalent to hardware reset and a noncommanded COMRESET may occur if there is an asynchronous loss of signal. Since COMRESET is equivalent to hardware reset, in the case of an asynchronous loss of signal some software settings may be lost without legacy mode software knowledge. In order to avoid losing important software settings without legacy mode driver knowledge, the software settings preservation ensures that the value of important software settings is maintained across a COMRESET. Software settings preservation may be enabled or disabled using SET FEATURES with a subcommand code of 06h (refer to section 13.3.6). If a device supports software settings preservation, the feature shall be enabled by default.

The software settings that shall be preserved across COMRESET are listed below. The device is only required to preserve the indicated software setting if it supports the particular feature/command the setting is associated with.

**INITIALIZE DEVICE PARAMETERS:** Device settings established with the INITIALIZE DEVICE PARAMETERS command. This command is obsolete in the ATA8-ACS standard, and was last defined in the ATA/ATAPI-5 standard.

**Power Management Feature Set Standby Timer:** The Standby timer used in the Power Management feature set.

**Read/Write Stream Error Log:** The Read Stream Error Log and Write Stream Error Logs (accessed using the GPL feature set commands).

**Security mode state:** The security mode state established by Security Mode feature set commands (refer to the ATA8-ACS standard). The device shall not transition to a different security mode state based on a COMRESET. For example, the device shall not transition from the SEC5: Unlocked / not Frozen state to state SEC4: Security enabled / Locked when a COMRESET occurs, instead the device shall remain in the SEC5: Unlocked / not Frozen state.

**SECURITY FREEZE LOCK:** The Frozen mode setting established by the SECURITY FREEZE LOCK command.

**SECURITY UNLOCK:** The unlock counter that is decremented as part of a failed SECURITY UNLOCK command attempt.

**SET MAX ADDRESS (EXT):** The maximum LBA specified in SET MAX ADDRESS or SET MAX ADDRESS EXT.

**[SET FEATURES \(Sense Data Reporting Enable/Disable\):](#)** [The enable/disable setting established by the SET FEATURES command with subcommand code of C3h.](#)

**SET FEATURES (Write Cache Enable/Disable):** The write cache enable/disable setting established by the SET FEATURES command with subcommand code of 02h or 82h.

**SET FEATURES (Set Transfer Mode):** PIO, Multiword, and UDMA transfer mode settings established by the SET FEATURES command with subcommand code of 03h.

**SET FEATURES (Advanced Power Management Enable/Disable):** The advanced power management enable/disable setting established by the SET FEATURES command with subcommand code of 05h or 85h. The advanced power management level established in the Count(7:0) register when advanced power management is enabled (SET FEATURES subcommand code 05h) shall also be preserved.

**SET FEATURES (Read Look-Ahead):** The read look-ahead enable/disable setting established by the SET FEATURES command with subcommand code of 55h or AAh.

**SET FEATURES (Release Interrupt):** The release interrupt enable/disable setting established by the SET FEATURES command with a subcommand code of 5Dh or DDh.

**SET FEATURES (SERVICE Interrupt):** The SERVICE interrupt enable/disable setting established by the SET FEATURES command with a subcommand code of 5Eh or DEh.

**SET FEATURES (Reverting to Defaults):** The reverting to power-on defaults enable/disable setting established by the SET FEATURES command with a subcommand code of CCh or 66h.

**SET MULTIPLE MODE:** The block size established with the SET MULTIPLE MODE command.

**NCQ QUEUE MANAGEMENT (Deadline Handling):** The state of WDNC and RDNC.

**Write-Read-Verify feature set:** The contents of IDENTIFY DEVICE data word 120 bit 1, words 210-211, and word 220 bits (7:0). The device shall not return to its Write-Read-Verify factory default setting after processing a COMRESET.

[editors note: update section 13.7.3 as noted below.]

### 13.7.3 Queued Error Log (10h)

The error-handling scheme for native queued commands halts processing of commands after the host is notified of an error on a native queued command. This allows host software to intervene and take appropriate action to resolve the error and avoids the potential for inconsistency due to data dependencies in the outstanding commands. The host explicitly restarts command processing by issuing a specific command to the device that results in the device aborting all remaining outstanding commands. Because the shadow Status and Error registers are not sufficiently large to contain both information about the error condition and the tag identifying the erring queued command, an additional log has been added in order for the host to be able to retrieve additional information for erring queued commands.

The General Purpose Logging (GPL) feature set is defined in [the ATA8-ACS ACS-2 standard](#).

If IDENTIFY DEVICE word 76 bit 15 is set to one, the Queued Error Log may be read using either of the READ LOG EXT or READ LOG DMA EXT commands.

If IDENTIFY DEVICE word 76 bit 15 is cleared to zero, the Queued Error Log shall be read using the READ LOG EXT command. An attempt to read the Queued Error Log using the READ LOG DMA EXT command shall be aborted and the state of the device shall not change.

Reading the Queued Error Log (10h) has the additional side effect defined in section [\[editors note: 13.6.2\]](#) of aborting any outstanding queued commands and returns a device that has halted due to a queued command error to a state where it has no commands outstanding and is again ready to accept commands (e.g., after completion of a command to read the log the device returns to state D10:Device\_idle state as defined in section [\[editors note: 11.2\]](#)). The Queued Error Log contains extended command error information.

The Queued Error Log reflects the error information for the first recorded NCQ command with error until such time as another NCQ error is encountered after reading the Queued Error Log. The contents of the Queued Error Log are indeterminate after a software reset or a COMRESET.

Devices supporting the native queued capability shall support the Queued Error Log. The Queued Error Log is one page in length and is defined in Figure 1.

[If the device supports NCQ Autosense \(i.e., IDENTIFY DEVICE word 78 bit 7 is set to one\), then:](#)

- a) [the Sense Key field;](#)
- b) [the Additional Sense Code field; and](#)
- c) [the Additional Sense Code Qualifier field](#)

[shall be set to values defined in SPC-4.](#)

[If the device does not support NCQ Autosense \(i.e., IDENTIFY DEVICE word 78 bit 7 is cleared to zero\), then:](#)

- d) [the Sense Key field shall be cleared to zero;](#)
- e) [the Additional Sense Code field shall be cleared to zero; and](#)
- f) [the Additional Sense Code Qualifier field shall be cleared to zero.](#)

Byte	7	6	5	4	3	2	1	0
0	NQ	UNL	R	TAG				
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(39:24)							
10	LBA(47:40)							
11	Reserved							
12	Count(7:0)							
13	Count(15:8)							
14	Reserved Sense Key							
15	Reserved Additional Sense Code							
16	Reserved Additional Sense Code Qualifier							
17	Reserved							
18	Reserved							
19	Reserved							
20	Reserved							
...								
255								
256								
...	Vendor Specific							
510								
511								

**Figure 1 – [Figure 221] Queued Error Log data structure definition**

- TAG      If the NQ bit is cleared to zero, the TAG field contains the TAG corresponding to the queued command that failed.
- UNL      If set 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 reading the Queued Error Log.

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(7:0) fields (bytes 3-5) in the log 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.

- NQ If set to one indicates that the error condition was a result of a non-queued command having been issued and that the TAG field is therefore not valid. If cleared to zero indicates that the TAG field is valid and that the error condition applies to a queued command.
- BYTE1-19 An image of a device to host Register FIS is embedded in the data structure. The fields correspond to the Shadow Register Block Registers and are encoded with error information consistent with the READ DMA QUEUED EXT or WRITE DMA QUEUED EXT command defined in the ATA8-ACS standard.
- ERROR The value corresponding to the ATA ERROR register value for the command that failed. The command-specific error condition of invalid tag value shall be handled as an invalid command parameter and shall be reported as such (i.e. ABRT bit set to one in the error register and all other bits cleared to zero).
- Note that the value returned in the ERROR field of the data structure is separate from the value returned in the Error shadow register when the initial error condition is signaled. The Error shadow register value is used for the purpose of signaling a queued command error, while the value in the ERROR field of the data structure provides specific information about the error condition that the specific queued command encountered.

[Sense Key](#)

[See SPC-4.](#)

[Additional Sense Code](#)

[See SPC-4](#)

[Additional Sense Code Qualifier](#)

[See SPC-4](#)

Vendor Specific

Allocated for vendor specific use.

Data Structure Checksum

The data structure checksum is the 2's complement of the sum of the first 511 bytes in the data structure. Each byte shall be added with unsigned arithmetic and overflow shall be ignored. The sum of all 512 bytes of the data structure is zero when the checksum is correct.

Reserved/R

All reserved fields shall be cleared to zero.