Proposed Draft **Serial ATA International Organization**

Version 14

October 29, 2012

TP_042v14_SATA31_Hybrid Information

Title: Hybrid Information Feature

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.

Editors note: please ignore "**Error! Reference source not found.**" in this revision. These are references within the original SATA Revision 3.1 document. The editor will address these on incorporation.

[editors note: Due to unexplained reasons, MANY MANY of the subheading numbers are very badly messed up. On incorporation, DO NOT DIRECTLY CUT/PASTE this text into your draft. Cleanse it by converting to plaintext (e.g. NOTEPAD) first.]

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 estopple 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.

Author Information

Author Name	Company	e-mail address
Jim Hatfield	Seagate	james.c.hatfield@seagate.com
James Boyd	Intel	james.a.boyd@intel.com
Curtis Stevens	Western Digital Corp	curtis.stevens@wdc.com

Workgroup Chair Information

Workgroup	Chairperson Name	e-mail address
Digital	Jim Hatfield	james.c.hatfield@seagate.com

Document History

Version	Date	Comments
00	11/28/2011	Initial draft
01	12/09/2011	1. Review comments from 11/28/11 and 12/05/11 2. Synchronize definition of 'medium' with ACS-2 3. Change LBA fields in NCQ NON-DATA to 'subcommand specific' 4. Change subcommand assignments to TBD 5. Explicitly say 'non-volatile caching medium' 6. Clarify PUIS requirements 7. Change 'range length' to 'sector count' for the two HYBRID CHANGE by xxx commands 8. Clarify sector count requirements if not enough data is in non-volatile caching medium 9. Clarify that the device cannot evict higher priority data if there are spinup restrictions 10. Misc minor editorial changes (spelling, etc.) (see additional red trikethrough text)
02	01/18/2012	1. Require support for NCQ Autosense 2. Specified SCSI sense codes for some errors 3. State that DevSleep 'should' be supported 4. State that Queuing Power Management 'shall' be supported 5. HYBRID DEMOTE BY SIZE a. Sector count specifies the 'minimum' number of sectors to

		be changed 6. HYBRID CHANGE BY LBA RANGE a. If device does not have 'sector count' logical blocks in the range for the stated priority, then return 'insufficient resources' error. 7. Hybrid Information Log a. Remove 'optimal write alignment' parameter 8. Auxiliary Field bits a. Delete the 'sequential access' bit 9. Interactions with power management a. (initial content has been added) 10. Misc spelling and punctuation changes 11. Changed DOWNLOAD MICROCODE text to an editor's note
3	Feb. 17, 2012	1. Incorporate comments from Jan. 20, 2012 2. Global: change 'media' to 'medium' 3. Global: change 'clean' to 'flush' 4. Change from 'shall support EPC' to 'should support EPC' 5. Added: 'should support APM' 6. HYBRID EVICT command a. if the Invalidate All field is set to one, then the device shall ignore all data transferred from the host 7. HYBRID DEMOTE BY SIZE a. The device shall complete the operation before returning command complete 8. Hybrid Information log a. Added field: Hybrid Health b. The Maximum Eviction Commands field of the log is now defined to be non-zero c. The offset of the first Hybrid Priority Descriptor structure was changed from 32 to 64 9. Hybrid Information feature a. Added subclause 'Additional Hybrid Errors' 10. Added additional DOWNLOAD MICROCODE questions to be answered
4	03/09/2012	Changes from 3/2/12 meeting of the authors: 1. Rearranged the 'shall', 'should' and 'may' support requirements for interdependent feature sets; 2. Globally changed 'flush' back to 'sync'; 3. Added text to the NCQ Send and Receive log to describe support

		for the HYBRID EVICT command 4. Resolved most of the TBDs and editors notes based on liaison responses from T13; 5. Removed all of the DOWNLOAD MICROCODE interactions, based on liaison with T13; 6. Accepted changes that were accepted in previous reviews; 7. Create a new option flag in the Hybrid Information log that determines how 'max priority' behaves. This bit is not changeable by the host. 8. If the bit is set and host cmd is read/write to max hybrid priority, then the device: a. SHALL insert data into NVM; b. SHALL flag an error when attempting to write and NVM is full; c. SHALL support the HYBRID EVICT command d. SHALL support the HYBRID CHANGE BY LBA command 9. If the bit is cleared and host cmd is read/write to max hybrid priority,, then: a. SHOULD insert data into NVM; b. SHALL NOT flag an error when attempting to write and NVM is full, but shall do VS eviction to make room for the new data; c. MAY support the HYBRID EVICT command d. MAY support the HYBRID CHANGE BY LBA command
5	03/16/2012	1. Incorporate changes suggested at Digital WG meeting of 3/12/2012 2. Incorporate changes suggested at authors meeting of 3/16/2012 3. Define 'evict' 4. Hybrid Information log a. GUID field: Change the field to a count of the number of times the feature has been enabled, and rename it to reflect that usage b. Hybrid Health field: define more precisely when each bit is set and cleared. (also subclause 13.20.10 Additional Hybrid Errors) c. Options field, MAXPRIORITY BEHAVIOR bit: more precisely define the meaning and usage of the bit 5. Change 'the data' to 'the requested data' in most cases 6. HYBRID CONTROL command a. If the DisableCachingMediafield is set to one, the device

		shall now ‘evict’ instead of ‘sync’ all data in the non-volatile caching medium
6	03/22/2012	1. Incorporate changes suggested at Digital WG meeting of 3/19/2012 2. Fix some incorrect IDENTIFY DEVICE word references (word 77->78 and word 78->79 in a few places) 3. Change the definition of ‘evict’ to remove the data, but not to sync the data 4. Split the description of the new Hybrid Health field a. Use ‘if the xxx bit is set/cleared, then…’ terminology here to describe the MEANING of the bits b. Move the requirements for SETTING the bits to a new section ‘Other Hybrid Errors’ in the feature set description in 13.x
7	03/26/2012	1. Incorporate changes suggested at Digital WG meeting of 3/26/12, as recorded in posted file: SATA31_TPR_D145_20120322_V06_Hybrid_Information_20120326.pdf
8	04/16/2012	1. Incorporate changes suggested at Digital WG meeting of 3/26/12, 04/02/12, 04/09/12, 04/16/12 as recorded in posted file: SATA31_TPR_D145_20120326_V07_Hybrid_Information – with KTM comments_20120417.pdf 2. Most of the comments were editorial (definitions, lists, punctuation, grammar, etc.) 3. HYBRID CONTROL command a. Disable Secondary Medium: changed from disabling both host and device usage, to only prohibiting use to store user data
9	05/01/2012	1. Misc editorial changes approved 04/30/2012: a. SATA31_TPR_D145_20120326_V07_Hybrid_Information – with KTMcomments_20120430.pdf b. SATA31_TPR_D145_20120419_V08_Hybrid_Information_clean_20120430.pdf
10	05/07/2012	1. Editor reviewed document and cleared TBDs. Updated section header references.
11	05/11/2012	1. Enabled comments and final format.
12	08/21/2012	1. Incorporate approved member review comment resolutions from files: a. TP_042v11_SATA31_Hybrid_Information_LBcomments_20120820.pdf b. TP_042v11_SATA31_Hybrid_Information_LBcomments_20120813_newman.pdf

13	10/20/2012	Incorporate approved member review comment resolutions from file TP_042v12_SATA31_Hybrid_InformationHJN+KTM_20121020.pdf
14	10/29/2012	Incorporate approved editorial comments from file: TP_042v13_SATA31_Hybrid_Information_20121029.pdf

Introduction

This describes a mechanism for the host to pass hints to a storage device about the data that is being transferred, when the storage device implements more than one storage medium.

Summary of the problem

Hybrid Hints Overview Section

Solid state drives (SSDs) provide excellent performance, power, form factor and ruggedness attributes. However, SSDs are significantly more expensive on a $/GB metric than a hard drive. Caching is intended to bring the benefits of increased performance and lower power to the mainstream. Caching solutions typically consist of host software that determines whether to place data on the hard drive or a small SSD acting as a cache. The Hybrid Information feature enables host software to provide information on data placement to a Solid State Hybrid Drive (SSHD), where the NVM is integrated with the hard drive itself. SSHDs have a smaller platform footprint than two separate devices (a hard drive and separate SSD), and also have a decreased cost since device resources (like CPU and DRAM) can be shared.

SSHDs determine data to cache based on observed workloads of LBAs and length of incoming requests. However, SSHDs do not possess host information to make the most optimal caching decisions (e.g., file type associated with the LBAs in the request). In order to make a more efficient SSHD the host should be involved in making caching decisions to aid in optimizing hit rate.

This proposal suggests a method for the host to indicate the caching priority of incoming requests to the device. The proposal also provides the host feedback on how much caching medium has been consumed at the various caching priority levels.

SSHDs may operate in one or more of these modes:

a) receive explicit requests from the host about what data to put in non-volatile caching medium (e.g. host-assisted);

b) observe data workloads without explicit host requests and decide internally what data to put in non-volatile caching medium (e.g. self-hinting); or

c) not use the non-volatile caching medium at all.

The highest priority passed to the device instructs the SSHD that this data shall be placed and remain in the non-volatile caching medium until explicitly evicted by the host. Intermediate caching priority levels inform the device of the importance of the data being placed in the non-volatile caching medium, but makes no requirement for the device to place the data in the non-volatile caching medium; the host only imparts a relative caching level in relation to other requests. The device should make the best decision possible based on the caching priority provided by the host and other device knowledge (e.g., rotational position optimizations).

An example would be the host issuing a request with a caching priority level of ‘MAX_PRIO’ to the device. In this case the device shall place the data in the non-volatile caching medium and the data shall remain there until the host evicts the data. Additional requests may be issued by the host with a caching level less than the ‘MAX_PRIO’ and this informs the SSHD of the priority of this data relative to other data at the same or lower caching priority levels. The device shall treat LBAs with a higher caching priority value as more suited for the non-volatile caching medium than LBAs with a lower caching priority value. This means that if the non-volatile caching medium on the SSHD has data in a lower caching priority (or the same caching priority) that this new data should take precedence for placement in the non-volatile caching medium. This may result in the device evicting data at a lower caching priority level to make room for the new higher priority data.

The lowest priority indicates that the device may place the data where it best deems fit (given spindle state, available caching medium capacity, etc).

Having multiple caching priorities allows for the host to group data by value. Data that is required to be in the non-volatile caching medium is passed with the highest priority level (e.g. connect standby data, boot data, etc). The intermediate priorities may be used for data that would provide value to the user if present in NVM, but is not required to meet power budgets or responsiveness criterion (e.g. medium files, application data, etc).

In order to maintain a given responsiveness for insertions into the non-volatile caching medium, the host may provide high and low dirty % thresholds to ensure that the device does not consume too much bandwidth syncing data between the non-volatile caching medium and primary medium, while still leaving room to absorb new writes.

A new log is added to provide the host with needed information on each caching priority level and contains information such as consumed caching medium for each caching priority level and the best write granularity for the non-volatile caching medium.

With these mechanisms, the host and device have a better opportunity to ensure hit rate and power requirements for SSHDs.

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.~~]

The changes proposed are:

a) add a few definitions;

b) defining additional command-dependent bits in the Register Host to Device FIS (27h) for use with the Hybrid Information feature;

c) renaming the NCQ QUEUE MANAGEMENT command to NCQ NON-DATA;

d) renaming the NCQ Queue Management log to NCQ NON-DATA log;

e) adding some new subcommands to SET FEATURES

   A) SATA-IO Enable/Disable Hybrid Information;

f) defining a new feature: Hybrid Information, which defines new optional behavior for these commands:

   A) NCQ NON-DATA;

   B) READ DMA EXT;

   C) WRITE DMA EXT;

   D) WRITE DMA FUA EXT;

   E) READ FPDMA QUEUED; and

   F) WRITE FPDMA QUEUED;

g) defining supported bits in IDENTIFY DEVICE and the Identify Device Data log;

h) adding a new SATA-only log:

   A) Hybrid Information ;

i) adding new subcommands to SEND FPDMA QUEUED command
   A) HYBRID EVICT;
j) adding new subcommands to NCQ NON-DATA command
   A) HYBRID DEMOTE BY SIZE;
   B) HYBRID CHANGE BY LBA RANGE;
   C) HYBRID CONTROL;
   and
k) power management changes:
   A) Related T13 proposals:
      a) f11134r0-EPC_Low_Power_Standby;
      b) f11133r0-EPC_Power_Source;
   B) Related SATA-IO proposals:
      a) TPR_D146_Queing Power Management;
   C) ability to spinup without blocking new I/O:
      a) device should have SET FEATURES PUIS enabled;
      b) device shall not support SET FEATURES PUIS feature set device spin-up subcommand;
   D) ability to spin down without blocking new I/O:
      a) see TPR_D146_Queing Power Management;
   E) ability for host to indicate it is unsafe to spinup to satisfy requests:
      a) see f11134r0-EPC_Low_Power_Standby;
      and
   F) host request for voluntary spin down:
      a) see f11134r0-EPC_Low_Power_Standby.

13 Revision History

[editors note: make no changes to clause 1]

14 Scope

[editors note: make no changes to clause 2]

15 Normative References

[editors note: make no changes to clause 3]

4 Definitions, abbreviations and conventions

4.1 Terminology

4.1.5 Definitions and abbreviations

[editors note: add these definitions to clause 4 (in sorted order) as noted below]

4.1.1.a dirty data

Dirty data is user data in a caching medium that is newer than the corresponding data in the primary medium.

4.1.1.b primary medium

The primary medium is the medium to which all data is synchronized.

4.1.1.c caching medium

A caching medium is an optional medium that may contain a subset of user data from the primary medium. The caching medium may contain data that is newer or identical to the data on the primary medium.

4.1.1.d hybrid device

A hybrid device is a device that contains both a primary medium and a non-volatile caching medium.

4.1.1.e mapping resource

A mapping resource is a vendor specific mechanism that may be used by the device internally to describe the physical location and attributes of user data.

4.1.1.f sync

Sync is a process within the device where dirty data in the caching medium is copied to the primary medium.

4.1.1.g evict

Evict is a process within the device to remove data from the caching medium.

[editors note: please add this to the list of abbreviations and acronyms:

SSHD Solid State Hybrid Device (see 4.1.1.d)

5 General Overview

[editors note: make no changes to this clause]

6 Cables and Connectors

[editors note: make no changes to this clause]

7 Phy Layer

[editors note: make no changes to this clause]

8 OOB and Phy Power States

[editors note: make no changes to this clause]

9 Link Layer

[editors note: make no changes to this clause]

[editors note: make the following changes to clause 10]

10 [Editor's note 10]Transport Layer

10.1 [Editor's note 10.1]Overview

10.2 [Editor's note 10.2]Frame Information Structure (FIS)

10.3 [Editor's note 10.3]FIS Types

10.3.5 [Editor's note 10.3.1]FIS Types scope

10.3.6 ~~<10.3.1>~~[Editor's note 10.3.2]FIS Type values

10.3.6.4 [Editor's note 10.3.2.1]FIS Type values overview

10.3.7 ~~<10.3.2>~~[Editor's note 10.3.3]CRC Errors on Data FISes

10.3.8 ~~<10.3.3>~~[Editor's note 10.3.4]All FIS types

10.3.9 ~~<10.3.4>~~[Editor's note 10.3.5]Register Host to Device FIS

10.3.9.4 [Editor's note 10.3.5.1]Register Host to Device FIS layout

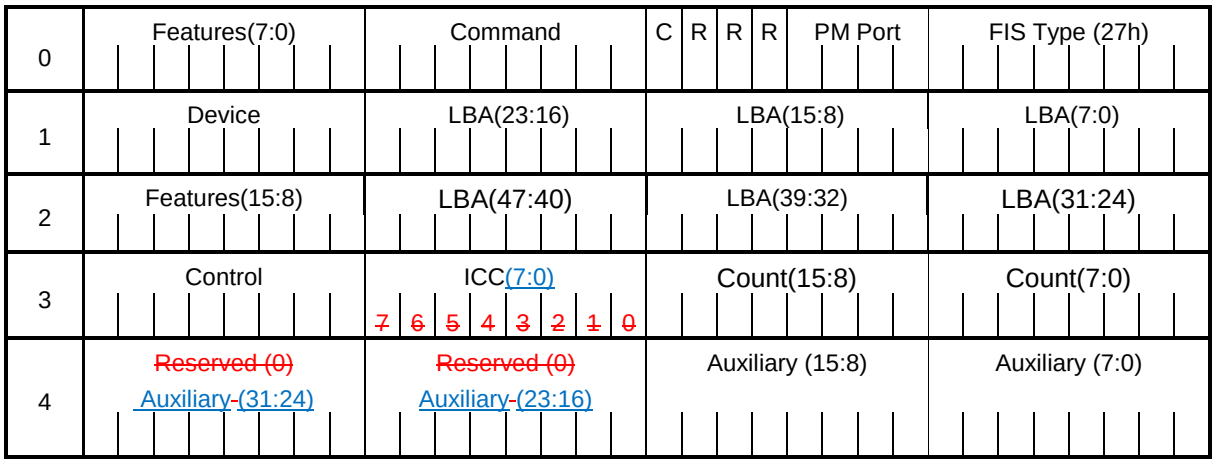

Figure 213 – [editors note: figure 213] Register - Host to Device FIS layout

If a field in this FIS is not defined by a command, it shall be Reserved for that command.

Field Definitions

FIS Type - Set to a value of 27h. Defines the rest of the FIS fields. Defines the length of the FIS as five Dwords.

C - This bit is set to one ~~when~~if the register transfer is due to an update of the Command register. The bit is cleared to zero ~~when~~if the register transfer is due to an update of the Device Control register. Setting C bit to one and SRST bit to one in the Device Control Field is invalid and results in indeterminate behavior.

Command - Contains the contents of the Command register of the Shadow Register Block.

Control - Contains the contents of the Device Control register of the Shadow Register Block.

LBA(7:0) - Contains the contents of the LBA ~~L~~low register of the Shadow Register Block.

~~Control - Contains the contents of the Device Control register of the Shadow Register Block.~~

LBA(15:8) - Contains the contents of the LBA ~~M~~mid register of the Shadow Register Block.

LBA(23:16) - Contains the contents of the LBA ~~H~~high register of the Shadow Register Block.

LBA(31:24) – Contains the contents of the expanded address field of the Shadow Register Block

LBA(39:32) – Contains the contents of the expanded address field of the Shadow Register Block

~~LBA(23:16) - Contains the contents of the LBA High register of the Shadow Register Block.~~

LBA(47:40) – Contains the contents of the expanded address field of the Shadow Register Block

Device - Contains the contents of the Device register of the Shadow Register Block.

Features(7:0) - Contains the contents of the Features register of the Shadow Register Block.

Features(15:8) – Contains the contents of the expanded address field of the Shadow Register Block

PM Port – ~~When~~If an endpoint device is attached via a Port Multiplier, specifies the device port address that the FIS should be delivered to. This field is set by the host.

R – Reserved – shall be cleared to zero.

Count(7:0) - Contains the contents of the Sector Count register of the Shadow Register Block.

Count(15:8) – Contains the contents of the expanded address field of the Shadow Register Block

~~LBA(7:0) - Contains the contents of the LBA Low register of the Shadow Register Block.~~

~~LBA(31:24) – Contains the contents of the expanded address field of the Shadow Register Block~~

ICC(7:0) - Isochronous Command Completion (ICC) contains a value is set by the host to inform device of a time limit. If a command does not define the use of this field, it shall be reserved.

Auxiliary-(7:0) – Contains parameter values specified on a per command basis.

Auxiliary-(15:8) – Contains parameter values specified on a per command basis.

Auxiliary-(23:16) – Contains parameter values specified on a per command basis.

Auxiliary-(31:24) – Contains parameter values specified on a per command basis.

10.3.4.1 ~~<10.3.4.1>~~[Editor's note 10.3.5.2]Description

The Register – Host to Device FIS is used to transfer the contents of the Shadow Register Block from the host to the device. This is the mechanism for issuing ATA commands to the device.

10.3.4.2 ~~<10.3.4.2>~~[Editor's note 10.3.5.3]Transmission

Transmission of a Register – Host to Device FIS is initiated by a write operation to either the command register, or a write to the Device Control register with a value different than is currently in the Device Control register in the host adapter's Shadow Register Block. Upon initiating transmission, the current contents of the Shadow Register Block are transmitted and the C bit in the FIS is set according to whether the transmission was a result of the Command register being written or the Device Control register being written. The host adapter shall set the BSY bit in the shadow Status register to one within 400 ns of the write operation to the Command register that initiated the transmission. The host adapter shall set the BSY bit in the shadow Status register to one within 400 ns of a write operation to the Device Control register if the write to the Device Control register changes the state of the SRST bit from zero to one. The host adapter shall not set the BSY bit in the shadow Status register for writes to the Device Control register that do not change the state of the SRST bit from zero to one.

NOTE 42 -It is important to note that Serial ATA host adapters enforce the same access control to the Shadow Register Block as parallel ATA devices enforce to the Command Block Registers.

Specifically, the host ~~is prohibited~~shall not write ~~from~~ ~~writing~~ the Features(7:0), Count(7:0), LBA(7:0), LBA(15:8), LBA(23:16), or Device registers ~~when~~if either BSY bit or DRQ bit is set to one in the Status Register. Any write to the Command Register ~~when~~if BSY bit or DRQ bit is set is ignored unless the write is to issue a Device Reset command.

[Editor's note: when updating SATA Revision 3.next, please do a global scan for bits 'set in' a register and change (appropriately) to 'bit set to one in' or 'bit cleared to zero in'.]

10.3.4.3 ~~<10.3.4.3>~~[Editor's note 10.3.5.4]Reception

Upon reception of a valid Register - Host to Device FIS the device updates its local copy of the Command and Control Block Register contents. Then the device either initiates execution of the command indicated in the Command register or initiates execution of the control request indicated in the Device Control register, depending on the state of the C bit in the FIS.

There are legacy BIOS and drivers that write the Device Control register to enable the interrupt just prior to issuing a command. To avoid unnecessary overhead, this FIS is transmitted to the device only upon a change of state from the previous value.

[editors note: make no changes to clause 11]

11 [Editor's note 11]Device Command Layer Protocol

[editors note: make no changes to clause 12]

12 [Editor’s note 12]Host Command Layer Protocol

[editors note: make the following changes to clause 13]

13 [Editor's note 13]Application Layer

13.1 [Editor's note 13.1]Parallel ATA Emulation

13.1.5 [Editor's note 13.1.1]Parallel ATA Emulation overview

[editors note: make no changes]

13.2 [Editor's note 13.2]IDENTIFY (PACKET) DEVICE

13.2.5 [Editor's note 13.2.1]IDENTIFY (PACKET) DEVICE overview

13.2.6 ~~<13.2.1>~~[Editor's note 13.2.2]IDENTIFY DEVICE

13.2.6.4 [Editor's note 13.2.2.1]IDENTIFY DEVICE information

[editors note: make the following changes to IDENTIFY DEVICE]

Table 84 – [editors note: table 84] IDENTIFY DEVICE information

Word	O/M	F/V	Bit	
78	O	F	9	~~Reserved~~ Hybrid Information feature is supported [Editors note: See ECN D158, TPR 038, and TPR 040]
79	O	V	9	~~Reserved~~ Hybrid Information feature is enabled [Editors note: See ECN D158, TPR 038, and TPR 040]

13.2.1.14 ~~<13.2.1.1>~~[Editor's note 13.2.2.2]Word 0~~ – ~~..46~~:~~

Set as indicated in ~~ATA8-ACS~~ACS-2

13.2.1.14 ~~<13.2.1.18>~~[Editor's note 13.2.2.19]Word 78: Serial ATA features supported

Bit 9, if set to one, then the device supports the Hybrid Information feature (see [editors note: 13.20]). If the device does not support NCQ (i.e., IDENTIFY DEVICE word 76 bit 8 is cleared to zero), then bit 9 shall be cleared to zero.

13.2.1.14 ~~<13.2.1.19>~~[Editor's note 13.2.2.20]Word 79: Serial ATA features enabled

Bit 9, if set to one, then the Hybrid Information feature is enabled. If the device does not support the Hybrid Information feature (i.e., IDENTIFY DEVICE word 78 bit 9 is cleared to zero), then word 79 bit 9 shall be cleared to zero.

13.2.7 ~~<13.2.2>~~[Editor's note 13.2.3]IDENTIFY PACKET DEVICE

13.2.7.4 [Editor's note 13.2.3.1]IDENTIFY PACKET DEVICE

[editors note: make no changes to IDENTIFY PACKET DEVICE]

13.2.8 ~~<13.2.3>~~[Editor's note 13.2.4]Determining Support for Serial ATA Features

[editors note: make no changes]

13.3 [Editor's note 13.3] SET FEATURES

13.3.5 [Editor's note 13.3.1] SET FEATURES overview

[editors note: make the following changes to SET FEATURES]

Devices are informed of host capabilities and have optional features enabled/disabled through the SET FEATURES command defined in the ATA8-ACS standard. Serial ATA features are controlled using a features value as defined in Table 87.

Table 87 – Features enable/disable values

Features(7:0) Value	Description
10h	Enable use of Serial ATA feature
90h	Disable use of Serial ATA feature

Count(7:0) contains the specific Serial ATA feature to enable or disable. The specific Serial ATA features in which SET FEATURES is applicable are defined in Table 88.

Table 88 – Feature identification values

Count(7:0) Value	Description
00h	Reserved
01h	Non-zero buffer offset in DMA Setup FIS
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	Device Automatic Partial to Slumber transitions
08h	Enable Hardware Feature Control
09h	Enable DevSleep
0Ah	Enable/Disable Hybrid Information
0Bh..~~09h -~~FFh	Reserved for future Serial ATA definition

13.3.6 ~~<13.3.1>~~[Editor's note 13.3.2] Enable/Disable Non-Zero Offsets in DMA Setup

[editors note: no changes]

13.3.7 ~~<13.3.2>~~[Editor's note 13.3.3]Enable/Disable DMA Setup FIS Auto-Activate Optimization

[editors note: no changes]

13.3.8 ~~<13.3.3>~~[Editor's note 13.3.4]Enable/Disable Device-Initiated Interface Power State Transitions

[editors note: no changes]

13.3.9 ~~<13.3.4>~~[Editor's note 13.3.5]Enable/Disable Guaranteed in-Order Data Delivery

[editors note: no changes]

13.3.10 ~~<13.3.5>~~[Editor's note 13.3.6]Enable/Disable Asynchronous Notification

[editors note: no changes]

13.3.11 ~~<13.3.6>~~[Editor's note 13.3.7]Enable/Disable Software Settings Preservation

[editors note: no changes]

13.3.12 ~~<13.3.7>~~[Editor's note 13.3.8]Enable/Disable Device Automatic Partial to Slumber Transitions

[editors note: no changes]

13.3.13 ~~<13.3.8>~~[Editor's note 13.3.9]Enable Hardware Feature Control

[editors note: no changes]

13.3.14 [Editor's note 13.3.10]Device Sleep

[editors note: TPR038 added Device Sleep]

[editors note: add this new section: Enable/Disable Hybrid Information]

13.3.15 [Editor's note 13.3.11]Enable/Disable Hybrid Information

See 13.20 for additional information about the Hybrid Information feature.

The Enable/Disable Hybrid Information subcommand:

a) enables the Hybrid Information feature and the non-volatile caching medium; or

b) disables the Hybrid Information feature and leaves the non-volatile caching medium in a vendor specific state.

The device shall return command aborted if the Hybrid Information feature is not supported.

If the Hybrid Information feature is enabled, then it shall remain enabled across all resets (e.g., power cycles), except as specified in 13.20.10.

13.3.15.4 Enable Hybrid Information feature subcommand

The Enable Hybrid Information Feature subcommand enables the Hybrid Information feature.

If the Hybrid Information feature is currently enabled (i.e., the Enabled field in the Hybrid Information log is set to FFh), then the device shall return command aborted.

If the Hybrid Information feature is currently disabled, then the device shall:

a) enable the Hybrid Information feature (i.e., set IDENTIFY DEVICE data word 79 bit 9 to one);

b) set the Enabled field (see 13.7.5.4.3) in the Hybrid Information log to FFh;

c) increment the Enable Count field (see 13.7.5.4.14) by one in the Hybrid Information log; and

d) enable the use of the non-volatile caching medium.

13.3.15.5 Disable Hybrid Information subcommand

The Disable Hybrid Information Feature subcommand disables the Hybrid Information feature.

If the Hybrid Information feature is currently enabled (i.e., the Enabled field in the Hybrid Information log is set to FFh), then the device shall:

a) disable the Hybrid Information feature (i.e., clear IDENTIFY DEVICE data word 79 bit 9 to zero);

b) set the Enabled field (see 13.7.5.4.3) in the Hybrid Information log to 00h; and

c) change the Hybrid Priority for all logical sectors in the non-volatile caching medium to zero.

If the Hybrid Information feature is currently disabled, then the device should return command completed with no error.

13.4 [Editor's note 13.4]Device Configuration Overlay

[editors note: make no changes]

13.5 [Editor's note 13.5]Software Settings Preservation (Optional)

[editors note: make no changes]

13.6 [Editor's note 13.6]Native Command Queueing (Optional)

13.6.5 [Editor's note 13.6.1]Native Command Queueing overview (Optional)

13.6.6 ~~<13.6.1>~~[Editor's note 13.6.2]Definition

[editors note: make no changes]

13.6.7 ~~<13.6.2>~~[Editor's note 13.6.3]Intermixing Non-Native Queued Commands and Native Queued Commands

[editors note: make no changes]

13.6.8 ~~<13.6.3>~~[Editor's note 13.6.4]READ FPDMA QUEUED

13.6.8.4 [Editor's note 13.6.4.1]READ FPDMA QUEUED command definition

[editors note: make the following changes to READ FPDMA QUEUED]

Queued native read commands make use of a command. The command supports LBA mode only and uses 48-bit addressing only. The format of the command is defined in Figure 224.

[editors note: the LBA fields have been re-sorted in bit order]

~~Register~~ Field	7	6	5	4	3	2	1	0
Features(7:0)	Sector Count 7:0							
Features(15:8)	Sector Count 15:8							
Count(7:0)	TAG					Reserved		
Count(15:8)	PRIO(1:0)		Reserved					
LBA(7:0)	LBA-(7:0)							
LBA(15:8)	LBA-(15:8)							
LBA(23:16)	LBA-(23:16)							
LBA(31:24)	LBA-(31:24)							
LBA(39:32)	LBA-(39:32)							
LBA(47:40)	LBA-(47:40)							
ICC(7:0)	ICC(7:0)							
Auxiliary(7:0)	Reserved							
Auxiliary-(15:8)	Reserved							
Auxiliary-(23:16)	Hybrid Information							
Auxiliary-(31:24)	Reserved							
Device	FUA	1	Res	0	Reserved			
Command	**60h**							

Figure 224 – READ FPDMA QUEUED command definition

TAG The TAG value shall be assigned by host software to be different from all other TAG values corresponding to outstanding commands. The assigned TAG value shall not exceed the value specified in IDENTIFY DEVICE word 75.

PRIO The Priority (PRIO) value is assigned by the host based on the priority of the command issued. The device should complete high priority requests in a more

timely fashion than normal and isochronous requests. The device should complete isochronous requests prior to its associated deadline.

00b	Normal Priority
01b	Isochronous – deadline dependent priority
10b	High priority
11b	Reserved

ICC — The Isochronous Command Completion (ICC) field is valid ~~when~~if PRIO is set to a value of 01b. It is assigned by the host based on the intended deadline associated with the command issued. ~~When~~If a deadline has expired, the device shall continue to complete the command as soon as possible. This behavior may be modified by the host if the device supports the NCQ NON-DATA ~~QUEUE MANAGEMENT~~ command (see 13.6.11) and supports the ~~Deadline Handling~~ DEADLINE HANDLING subcommand (see [editors note: 13.6.5.1xxx]). This subcommand allows the host to set whether the device shall abort (or continue processing) commands that have exceeded the time set in ICC.

There are several parameters encoded in the ICC field: Fine or Coarse timing, Interval and the Max Time. The Interval indicates the time units of the Time Limit parameter.

If ICC Bit 7 is cleared to zero, then:

- a) ~~T~~the time interval is fine-grained;~~.~~
- b) ~~I~~interval = 10 ms~~ec~~;
- c) ~~T~~time ~~L~~limit = (ICC[6:0] + 1) ~~*~~× 10 ms~~ec~~; and
- d) ~~M~~max ~~F~~fine ~~T~~time = 128 ~~*~~× 10 ms~~ec~~ = 1.28 s~~ec~~.

If ICC Bit 7 is set to one (coarse encoding), then:

- a) ~~T~~the time interval is coarse-grained;
- b) ~~I~~interval = 0.5 s~~ec~~;
- c) ~~T~~time ~~L~~limit = (ICC[6:0] + 1) ~~*~~× 0.5 s~~ec~~; and
- d) ~~M~~max ~~C~~coarse ~~T~~time = 128 ~~*~~× 0.5 s~~ec~~ = 64 s~~ec~~.

FUA — ~~When~~If set to one forces ~~the data~~the requested data to be retrieved from non-volatile ~~the~~ storage medium regardless of whether the storage device holds the requested information in its ~~buffers or~~volatile cache. If the device holds a modified copy of the requested data as a result of having cached writes, the modified data is first written to the medium before being retrieved from the storage medium as part of this operation. ~~When~~If cleared to zero ~~the data~~the requested data may be retrieved either from the device's storage non-volatile storage medium or from ~~buffers~~volatile ~~/~~cache that the device may include.

Hybrid Information — If the Hybrid Information feature is supported (i.e., IDENTIFY DEVICE data word 78 bit 9 is set to one) and is enabled (i.e., IDENTIFY DEVICE data word 79 bit 9 is set to one), then the Hybrid Information field shall be processed as defined in 13.20. If the Hybrid Information feature is not supported (i.e., IDENTIFY DEVICE data word 78 bit 9 is cleared to zero) or is disabled (i.e., IDENTIFY DEVICE data word 79 bit 9 is cleared to zero) then the device shall ignore the Hybrid Information field.

Others — All other registers have contents consistent with the READ DMA QUEUED EXT command defined in the ATA8-ACS standard, including the Sector Count ~~(~~(15:0) convention where a value of zero specifies that 65~~,~~ 536 sectors are to be transferred.

[editors note: the rest of section 13.6.3 is unchanged]

13.6.9 ~~<13.6.4>~~[Editor's note 13.6.5]WRITE FPDMA QUEUED

13.6.10 [Editor's note 13.6.5.1]WRITE FPDMA QUEUED command definition

[editors note: make the following changes to WRITE FPDMA QUEUED]

Queued native write commands make use of a command. The format of the command is defined in Figure 229.

[editors note: the LBA fields have been re-sorted in bit order]

~~Register~~ Field	7	6	5	4	3	2	1	0
Features(7:0)	Sector Count 7:0							
Features(15:8)	Sector Count 15:8							
Count(7:0)	TAG					Reserved		
Count(15:8)	PRIO(1:0)		Reserved					
LBA(7:0)	LBA 7:0							
LBA(15:8)	LBA 15:8							
LBA(23:16)	LBA 23:16							
LBA(31:24)	LBA 31:24							
LBA(39:32)	LBA 39:32							
LBA(47:40)	LBA 47:40							
ICC(7:0)	ICC(7:0)							
Auxiliary(7:0)	Reserved							
Auxiliary (15:8)	Reserved							
Auxiliary (23:16)	Hybrid Information							
Auxiliary (31:24)	Reserved							
Device	FUA	1	0	0	Reserved			
Command	**61h**							

Figure 229 – WRITE FPDMA QUEUED command definition

TAG The TAG value shall be assigned by host software to be different from all other TAG values corresponding to outstanding commands. The assigned TAG value shall not exceed the value specified in IDENTIFY DEVICE word 75.

PRIO — The Priority (PRIO) value is assigned by the host based on the priority of the command issued. The device should complete high priority requests in a more timely fashion than normal and isochronous requests. The device should complete isochronous requests prior to its associated deadline.

00b	Normal Priority
01b	Isochronous – deadline dependent priority
10b	High priority
11b	Reserved

ICC — The Isochronous Command Completion (ICC) field is valid ~~when~~if PRIO is set to a value of 01b. It is assigned by the host based on the intended deadline associated with the command issued. ~~When~~If a deadline has expired, the device shall continue to complete the command as soon as possible. This behavior may be modified by the host if the device supports the NCQ NON-DATA ~~QUEUE MANAGEMENT~~ command (see 13.6.11) and supports the ~~Deadline Handling~~ DEADLINE HANDLING subcommand (see [editors note: 13.6.9.xxx]). This subcommand allows the host to set whether the device shall abort (or continue processing) commands that have exceeded the time set in ICC.

There are several parameters encoded in the ICC field:

a) ~~F~~fine or ~~C~~coarse timing, ~~I~~interval; and

b) the ~~M~~max ~~T~~time.

The Interval indicates the time units of the Time Limit parameter.

If ICC Bit 7 is cleared to zero, then:

- a) ~~T~~the time interval is fine-grained;~~.~~
- b) ~~I~~interval = 10 ms~~ec~~;
- c) ~~T~~time ~~L~~limit = (ICC[6:0] + 1) ~~*~~× 10 ms~~ec~~; and
- d) ~~M~~max ~~F~~fine ~~T~~time = 128 ~~*~~× 10 msec = 1.28 s~~ec~~.

If ICC Bit 7 is set to one (coarse encoding), then:

- a) ~~T~~the time interval is coarse-grained;
- b) ~~I~~interval = 0.5 s~~ec~~;
- c) ~~T~~time ~~L~~limit = (ICC[6:0] + 1) ~~*~~× 0.5 s~~ec~~; and
- d) ~~M~~max ~~C~~coarse ~~T~~time = 128 ~~*~~× 0.5 s~~ec~~ = 64 s~~ec~~.

FUA — ~~When~~If set to one forces ~~the data~~the requested data to be written to non-volatile ~~the~~ storage ~~media~~medium before completion status is indicated. When cleared to zero the device may indicate completion status before ~~the data~~the requested data is committed to the ~~media~~non-volatile storage medium.

Hybrid Information — If the Hybrid Information feature is supported (i.e., IDENTIFY DEVICE data word 78 bit 9 is set to one) and is enabled (i.e., IDENTIFY DEVICE data word 79 bit 9 is set to one), then the Hybrid Information field shall be processed as defined in 13.20). If the Hybrid Information feature is not supported (i.e., IDENTIFY DEVICE data word 78 bit 9 is cleared to zero) or is disabled (i.e., IDENTIFY DEVICE data word 79 bit 9 is cleared to zero) then the device shall ignore the Hybrid Information field.

Others — All other registers as specified for the WRITE DMA QUEUED EXT command defined in the ATA8-ACS standard, including the Sector Count ~~-~~(15:0) convention where a value of zero specifies that 65~~,~~536 sectors are to be transferred.

[editors note: the rest of section 13.6.4 is unchanged]

13.6.11 ~~<13.6.5>~~[Editor's note 13.6.6]NCQ NON-DATA ~~QUEUE MANAGEMENT~~

13.6.11.4 [Editor's note 13.6.6.1]NCQ NON-DATA command definition

[editors note: make the following changes to NCQ QUEUE MANAGEMENT]

[editors note: heading numbering is messed up for all of subclause 13.6.9]
The NCQ ~~Queue Management~~ NON-DATA feature allows the host to send non-data commands to the device without aborting the queue (see 13.6.3.3) ~~manage the outstanding NCQ commands and/or affect the processing of NCQ commands~~.

~~The NCQ QUEUE MANAGEMENT~~ NON-DATA ~~command is a non-data NCQ command.~~ Some ~~Only specified~~ NCQ ~~QUEUE MANAGEMENT~~ NON-DATA subcommands are ~~executed~~processed as Immediate NCQ commands (see 4.1.1.70xxx).

The NCQ NON-DATA log (see 13.7.9) indicates which subcommands are supported.

If NCQ is disabled and an NCQ ~~QUEUE MANAGEMENT~~ NON-DATA command is issued to the device, then the device shall abort the command with the ERR bit set to one in the Status register and the ABRT bit set to one in the Error register. This command ~~is prohibited~~shall not be implemented for devices that implement the PACKET feature set. The queuing behavior of the device depends on which subcommand is specified.

[editors note: the LBA fields have been re-sorted in bit order]

~~Register~~ Field	7	6	5	4	3	2	1	0
Features(7:0)	Subcommand Specific				Subcommand			
Features(15:8)	~~Reserved~~ Subcommand Specific							
Count(7:0)	TAG					Reserved		
Count(15:8)	~~Reserved~~ Subcommand Specific							
LBA(7:0)	Subcommand Specific ~~(TTAG)~~ [editors note: combine bits 7:0 as one field]					Subcommand Specific ~~Reserved~~		
LBA(15:8)	~~Reserved~~ Subcommand Specific							
LBA(23:16)	~~Reserved~~ Subcommand Specific							
LBA(31:24)	~~Reserved~~ Subcommand Specific							
LBA(39:32)	~~Reserved~~ Subcommand Specific							
LBA(47:40)	~~Reserved~~ Subcommand Specific							
ICC(7:0)	Reserved							
Auxiliary(7:0)	Reserved							
Auxiliary(15:8)	Reserved							
Auxiliary(23:16)	Subcommand Specific							
Auxiliary(31:24)	Reserved							

Device	Res	1	Res	0	Reserved
Command	**63h**				

Figure 233 – NCQ ~~QUEUE MANAGEMENT~~ NON-DATA - Command definition

TAG The TAG value shall be assigned by host software to be different from all other TAG values corresponding to outstanding commands. The assigned TAG value shall not exceed the value specified in IDENTIFY DEVICE word 75.

Subcommand The Subcommand field (see Table 89) defines the subcommands that are valid.

If an invalid subcommand is specified, then the device shall abort the command with the ERR bit set to one in the Status register, the ABRT bit set to one in the Error register, and shall cause all outstanding commands to be aborted.

Table 89 – Subcommand Field

Subcommand	Description	Reference
0h	Abort NCQ queue	13.6.11.5
1h	~~Deadline Handling~~ DEADLINE HANDLING	13.613.6.613.6.5.1
2h	HYBRID DEMOTE BY SIZE	11.313.6.313.6.5.4
3h	HYBRID CHANGE BY LBA RANGE	11.313.6.313.6.5.5
4h	HYBRID CONTROL	11.313.6.313.6.5.6
5~~2~~h ~~–~~ ..Fh	Reserved	

Subcommand Specific The Subcommand Specific fields are defined for each subcommand as referenced from Table 89.

In previous versions of this specification LBA(7:3) were assigned to Subcommand Specific (TTAG) and LBA(2:0) were reserved. See Table 90 and the associated subcommand definitions for Subcommand Specific (TTAG) bit mapping.

~~TAG The TAG value shall be assigned by host software to be different from all other TAG values corresponding to outstanding commands. The assigned TAG value shall not exceed the value specified in IDENTIFY DEVICE word 75.~~

~~Subcommand Specific (TTAG) is the selected queue TAG. This allows the host to select the specific outstanding queued command to be managed.~~

The error and normal returns for this command are subcommand specific.

13.6.11.5 ~~<13.6.5.1>~~[Editor's note 13.6.6.2]ABORT NCQ QUEUE Subcommand (0h)

13.6.5.4.1 [Editor's note 13.6.6.2.1] ABORT NCQ QUEUE Subcommand definition (0h)

A Subcommand set to 0h specifies the ABORT NCQ QUEUE subcommand. The ABORT NCQ QUEUE subcommand is an immediate NCQ command. Support for this subcommand is indicated in the NCQ NON-DATA ~~Queue Management~~ Log (see 13.7.9.5xxx).

The ABORT NCQ QUEUE subcommand shall affect only those NCQ commands for which the device has indicated command acceptance before accepting this NCQ NON-DATA ~~QUEUE MANAGEMENT~~ command.

The format of the command is defined in Figure 234.

[editors note: the LBA fields have been re-sorted in bit order]

~~Register~~ Field	7	6	5	4	3	2	1	0
Features(7:0)	Abort Type				0h			
Features(15:8)	Reserved							
Count(7:0)	TAG					Reserved		
Count(15:8)	Reserved							
LBA(7:0)	TTAG					Reserved		
LBA(15:8)	Reserved							
LBA(23:16)	Reserved							
LBA(31:24)	Reserved							
LBA(39:32)	Reserved							
LBA(47:40)	Reserved							
ICC(7:0)	Reserved							
Auxiliary(7:0)	Reserved							
Auxiliary(15:8)	Reserved							
Auxiliary(23:16)	Subcommand Specific							
Auxiliary(31:24)	Reserved							
Device	Res	1	Res	0	Reserved			
Command	**63h**							

Figure 234 – ABORT NCQ QUEUE ~~NCQ QUEUE MANAGEMENT, Abort NCQ Queue~~ - Command definition

Abort Type describes the action requested. Table 90 shows the defined abort types. The NCQ NON-DATA ~~Queue Management~~ Log (see 13.7.5) provides a list of abort types supported by the device.

TAG The TAG value shall be assigned by host software to be different from all other TAG values corresponding to outstanding commands. TAG shall not exceed the value specified in IDENTIFY DEVICE word 75.

TTAG The TTAG field contains the value of the TAG of the outstanding command that is requested to be aborted. The TTAG value is only valid ~~when~~if the Abort Type field is set to 3h (Abort Selected). TTAG shall not exceed the value specified in IDENTIFY DEVICE word 75.

Table 90 – Abort Type

Abort Type	Abort Type	Description
0h	Abort All	The device shall attempt to abort all outstanding NCQ commands.
1h	Abort Streaming	The device shall attempt to abort all outstanding NCQ Streaming commands. All non-streaming NCQ commands shall be unaffected.
2h	Abort Non-Streaming	The device shall attempt to abort all outstanding NCQ Non-Streaming commands. All NCQ Streaming commands shall be unaffected.
3h	Abort Selected	The device shall attempt to abort the outstanding NCQ command associated with the tag represented in TTAG field.
4h~~–~~..Fh		Reserved

13.6.5.4.2 ~~<13.6.5.1.1>~~[Editor's note 13.6.6.2.2]Success Outputs

If a supported Abort Type parameter is specified, then the device shall indicate success, even if the command results in no commands being aborted.

~~When~~If an Abort NCQ Queue command completes successfully, a Set Device Bits FIS shall be sent to the host to complete the Abort subcommand and commands that were aborted as a consequence of the Abort subcommand by setting the ACT bits for those commands to one. This SDB FIS may also indicate other completed commands.

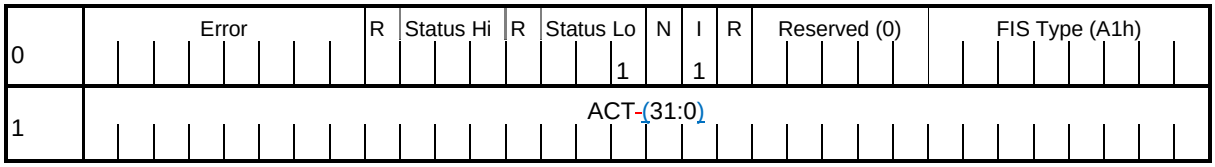

Figure 235 – ABORT NCQ QUEUE ~~NCQ QUEUE MANAGEMENT, Abort NCQ Queue~~ - Successful completion

ACT — The ACT field of the Set Device Bits FIS communicates completion notification for each of up to 32 commands. The field is bit-significant and the device sets bit positions to one for each command tag it is indicating completion notification for. The device may set more than one bit to one if it is explicitly aggregating successful status returns. The device shall set the appropriate bit to one for each queued command that has been aborted, and shall set to one the bit associated with the TAG value for the Abort NCQ Queue command.

Error — The Error register shall contain 00h.

Status — As defined in [Editor's note 10.3.6]. The ERR bit shall be cleared to zero indicating successful command completion. Bit 4 may be set to one.

I — Interrupt bit. The interrupt bit shall be set to one.

All other fields as defined in [Editor's note 10.3.6].

13.6.5.4.3 ~~<13.6.5.1.2>~~[Editor's note 13.6.6.2.3]Error Outputs

13.6.5.2.3.1 ~~<13.6.5.1.2.1>~~[Editor's note 13.6.6.2.3.1]Upon receipt of a command

If the value of the TTAG field equals the value of the TAG field, or if an unsupported Abort type parameter is specified, the device shall return command aborted.

If the device has received a command that has not yet been acknowledged by clearing the BSY bit to zero and an error is encountered, the device shall transmit a Register Device to Host FIS (see Figure 236) with the ERR bit set to one and the BSY bit cleared to zero in the Status field, the ATA error code in the Error field.

[editors note: the LBA fields have been re-sorted in bit order]

~~Register~~ Field	7	6	5	4	3	2	1	0
Error	ERROR							
Count(7:0)	na							
Count(15:8)	na							
LBA(7:0)	na							
LBA(15:8)	na							
LBA(23:16)	na							
LBA(31:24)	na							
LBA(39:32)	na							
LBA(47:40)	na							
Device	na							
Status	BSY	DRDY	DF	na	DRQ	na	na	ERR

Figure 236 – ABORT NCQ QUEUE ~~NCQ QUEUE MANAGEMENT, Abort NCQ Queue~~ - error on command receipt

ERROR	ATA error code for the failure condition of the failed command
BSY	0
DRDY	1
DF	0
DRQ	0
ERR	1

Status bit 4 may be set to one.

Following transmission of the Register Device to Host FIS, the device shall stop processing any outstanding or new commands until the Queued Error Log (see 13.7.7) has been read before continuing to abort all outstanding commands. See [Editor's note 13.6.3.3] for more details.

13.6.5.2.3.2 ~~<13.6.5.1.2.2>~~[Editor's note 13.6.6.2.3.2]During execution of a command

If all commands have been acknowledged by clearing the BSY bit to zero and an error condition is detected, the device shall transmit a Set Device Bits FIS (see [Editor's note Figure 232]) to the host with the ERR bit set to one in the Status field, the ATA error code in the Error field, and the Interrupt bit set to one. All outstanding commands at the time of an error are aborted as part of the error response and may be re-issued as appropriate by the host. For any commands that have not completed successfully or have resulted in error, the device shall clear the corresponding ACT bits to zero in the Set Device Bits FIS.

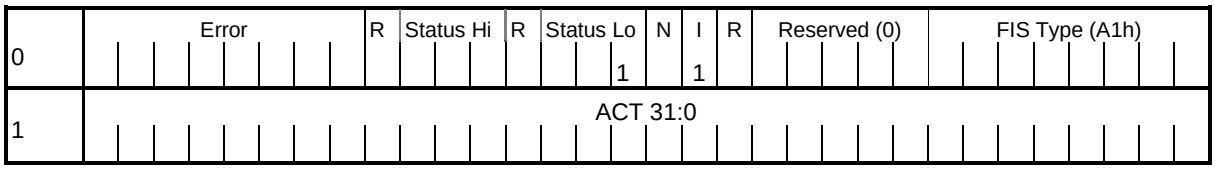

Figure 237 – ABORT NCQ QUEUE ~~NCQ QUEUE MANAGEMENT, Abort NCQ Queue~~ – error during execution

ACT	The ACT field of the Set Device Bits FIS communicates successful completion notification for each of up to 32 queued commands. The field is bit-significant and the device sets bit positions to one for each command tag it is indicating successful completion notification for. The device may set more than one bit to one if it is explicitly aggregating successful status returns.
Error	The Error register shall contain the ATA error code.
Status	As defined in [Editor's note 10.3.6]. The ERR bit shall be set to one indicating an NCQ error has occurred. Status bit 4 may be set to one.
I	Interrupt bit. The interrupt bit shall be set to one.

All other fields as defined in [Editor's note 10.3.6].

Only the registers that are updated as part of the Set Device Bits FIS are modified if the device signals an error condition ~~when~~if the BSY bit in the shadow Status register is cleared to zero, leaving the other Shadow Register Block Registers unchanged. If the device signals an error condition ~~when~~if the BSY bit in the shadow Status register is set to one, the device clears the BSY bit to zero with a Register Device to Host FIS which updates all registers in the Shadow Register Block.

Following transmission of the Set Device Bits FIS, the device shall stop processing any outstanding or new commands until the Queued Error Log (see [Editor's note 13.7.7]) has been read before continuing to abort all outstanding commands. See [Editor's note 13.6.3.3] for more details.

13.6.5.1 ~~<13.6.5.1.3>~~[Editor's note 13.6.6.3]DEADLINE HANDLING ~~Deadline Handling~~ Subcommand (1h)

[editors note: all of this section is renumbered to match the indentation level of the Abort NCQ command]

13.6.5.4.4 [Editor's note 13.6.6.3.1]DEADLINE HANDLING Subcommand definition (1h)

A Subcommand set to 1h specifies the ~~Deadline Handling~~ DEADLINE HANDLING Subcommand. This subcommand controls how NCQ Streaming commands are processed by the device. Support for this subcommand is indicated in the NCQ NON-DATA ~~Queue Management~~ Log (see 13.7.9.10xxx).~~Support for this subcommand is indicated in the NCQ Queue Management Log (see 13.7.5).~~ The format of the command is defined in Figure 238.

[editors note: the LBA fields have been re-sorted in bit order]

~~Register~~ Field	7	6	5	4	3	2	1	0
Features(7:0)	Reserved		RDNC	WDNC	1h			
Features(15:8)	Reserved							
Count(7:0)	TAG					Reserved		
Count(15:8)	Reserved							
LBA(7:0)	Reserved							
LBA(15:8)	Reserved							
LBA(23:16)	Reserved							
LBA(31:24)	Reserved							
LBA(39:32)	Reserved							
LBA(47:40)	Reserved							
ICC(7:0)	Reserved							
Auxiliary(7:0)	Reserved							
Auxiliary(15:8)	Reserved							
Auxiliary(23:16)	Subcommand Specific							
Auxiliary(31:24)	Reserved							
Device	Res	1	Res	0	Reserved			
Command	**63h**							

Figure 238 – DEADLINE HANDLING ~~NCQ QUEUE MANAGEMENT, Deadline Handling~~ - Command definition

WDNC — If the WDNC (Write Data Not Contiunue) bit is cleared to zero, then the device may allow WRITE FPDMA QUEUED command completion times to exceed what the ICC parameter specified. If the WDNC bit is set to one, then the all WRITE FPDMA QUEUED commands shall be completed by the time specified by the ICC timer value, otherwise the device shall return command aborted for all outstanding commands. WDNC is only applicable to WRITE FPDMA QUEUED

commands with PRIO is set to 01b (Isochronous – deadline dependent priority). (See [Editor's note 13.6.4]).

RDNC — If the RDNC (Read Data Not Continue) bit is cleared to zero, then the device may allow READ FPDMA QUEUED command completion times to exceed what the ICC parameter specified. If the RDNC bit is set to one, then all READ FPDMA QUEUED commands shall be completed by the time specified by the ICC timer value, otherwise the device shall return command aborted for all outstanding commands. RDNC is only applicable to READ FPDMA QUEUED commands with PRIO is set to 01b (Isochronous – deadline dependent priority). (See [Editor's note 13.6.3]).

The state of the WDNC and RDNC bits shall be preserved across software resets and COMRESETs (via Software Setting Preservations), and shall not be preserved across power cycles.

TAG — The TAG value shall be assigned by host software to be different from all other TAG values corresponding to outstanding commands. TAG shall not exceed the value specified in IDENTIFY DEVICE word 75.

13.6.5.4.5 ~~<13.6.5.1.3.1>~~[Editor's note 13.6.6.3.2]Success Outputs

If this ~~Deadline Handling~~ DEADLINE HANDLING Subcommand command is supported, the device shall return command completed with no error.

~~When~~If a~~n~~ ~~Deadline Handling~~ DEADLINE HANDLING Subcommand command completes successfully, a Set Device Bits FIS shall be sent to the host to complete the ~~Deadline Handling~~ DEADLINE HANDLING subcommand. This SDB FIS may also indicate other completed commands.

0	Error			R	Status Hi			R	Status Lo			N	I	R	Reserved (0)	FIS Type (A1h)
	7	6	5	4	3	2	1	0	7	6	5	4	1	0		
1	ACT (31:0)															
	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16
	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0

Figure 239 – DEADLINE HANDLING ~~NCQ QUEUE MANAGEMENT, Deadline Handling~~ - Successful completion

ACT — The ACT field of the Set Device Bits FIS communicates completion notification for each of up to 32 commands. The field is bit-significant and the device sets bit positions to one for each command tag it is indicating completion notification for. The device may set more than one bit to one if it is explicitly aggregating successful status returns. The device shall set to one the bit associated with the TAG value for the ~~Deadline Handling~~ DEADLINE HANDLING command.

Error — The Error register shall contain 00h.

Status — As defined in [Editor's note 10.3.6]. The ERR bit shall be cleared to zero indicating successful command completion. Bit 4 may be set to one.

I — Interrupt bit. The interrupt bit shall be set to one.

All other fields as defined in [Editor's note 10.3.6].

13.6.5.4.6 ~~<13.6.5.1.3.2>~~[Editor's note 13.6.6.3.3]Error Outputs

13.6.5.3.3.1 ~~<13.6.5.1.3.2.1>~~[Editor's note 13.6.6.3.3.1]Upon receipt of a command

If the device has received a command that has not yet been acknowledged by clearing the BSY bit to zero and an error is encountered, the device shall transmit a Register Device to Host FIS (see Figure 240) with the ERR bit set to one and the BSY bit cleared to zero in the Status field, the ATA error code in the Error field.

[editors note: the LBA fields have been re-sorted in bit order]

~~Register~~ Field	7	6	5	4	3	2	1	0
Error	ERROR							
Count(7:0)	na							
Count(15:8)	na							
LBA(7:0)	na							
LBA(15:8)	na							
LBA(23:16)	na							
LBA(31:24)	na							
LBA(39:32)	na							
LBA(47:40)	na							
Device	na							
Status	BSY	DRDY	DF	na	DRQ	na	na	ERR

Figure 240 – DEADLINE HANDLING ~~NCQ QUEUE MANAGEMENT, Deadline Handling~~ - error on command receipt

ERROR ATA error code for the failure condition of the failed command

BSY 0

DRDY 1

DF 0

DRQ 0

ERR 1

Status bit 4 may be set to one.

Following transmission of the Register Device to Host FIS, the device shall stop processing any outstanding or new commands until the Queued Error Log (see [Editor's note 13.7.7]) has been read before continuing to abort all outstanding commands. See [Editor's note 13.6.3.3] for more details.

13.6.5.3.3.2 ~~<13.6.5.1.3.2.2>~~[Editor's note 13.6.6.3.3.2]During execution of a command

If all commands have been acknowledged by clearing the BSY bit to zero and an error condition is detected, the device shall transmit a Set Device Bits FIS (see [Editor's note Figure 232]) to the host with the ERR bit set to one in the Status field, the ATA error code in the Error field, and the Interrupt bit set to one. All outstanding commands at the time of an error are aborted as part of

the error response and may be re-issued as appropriate by the host. For any commands that have not completed successfully or have resulted in error, the device shall clear the corresponding ACT bits to zero in the Set Device Bits FIS.

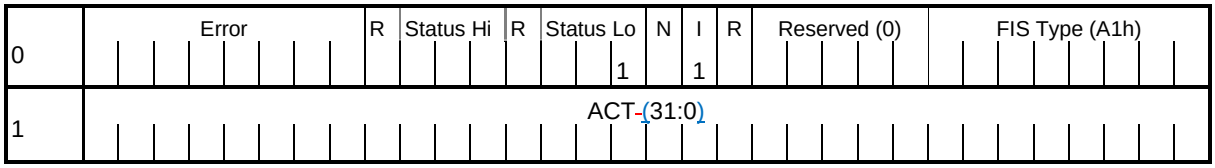

Figure 241 – DEADLINE HANDLING ~~NCQ QUEUE MANAGEMENT, Deadline Handling~~ – error during execution

ACT	The ACT field of the Set Device Bits FIS communicates successful completion notification for each of up to 32 queued commands. The field is bit-significant and the device sets bit positions to one for each command tag it is indicating successful completion notification for. The device may set more than one bit to one if it is explicitly aggregating successful status returns.
Error	The Error register shall contain the ATA error code.
Status	As defined in [Editor's note 10.3.6]. The ERR bit shall be set to one indicating an NCQ error has occurred. Status bit 4 may be set to one.
I	Interrupt bit. The interrupt bit shall be set to one.

All other fields as defined in [Editor's note 10.3.6].

Only the registers that are updated as part of the Set Device Bits FIS are modified if the device signals an error condition ~~when~~if the BSY bit in the shadow Status register is cleared to zero, leaving the other Shadow Register Block Registers unchanged. If the device signals an error condition ~~when~~if the BSY bit in the shadow Status register is set to one, the device clears the BSY bit to zero with a Register Device to Host FIS which updates all registers in the Shadow Register Block.

Following transmission of the Set Device Bits FIS, the device shall stop processing any outstanding or new commands until the Queued Error Log (see [Editor's note 13.7.7]) has been read before continuing to abort all outstanding commands. See [Editor's note 13.6.3.3] for more details.

[editors note: add the HYBRID DEMOTE BY SIZE subcommand to the NCQ NON-DATA command]

13.6.5.4 [Editor's note 13.6.6.4] HYBRID DEMOTE BY SIZE Subcommand (2h)

13.6.5.4.7 [Editor's note 13.6.6.4.1] HYBRID DEMOTE BY SIZE Subcommand definition

The HYBRID DEMOTE BY SIZE subcommand is used to change the hybrid priority associated with logical sectors. Support for this subcommand is indicated in the NCQ NON-DATA Log (see 13.7.9.13).

The device changes the Hybrid Priority of logical sectors in the non-volatile caching medium from the value specified in the ~~From~~ Priority field (see [Editor's note 13.20.2.2]Hybrid Priority) in the Hybrid Information field to the Hybrid Priority value specified in the Hybrid Priority field (see 13.20.6.5).

Table 91 specifies the number of logical sectors that should be demoted by the HYBRID DEMOTE BY SIZE. The number of logical sectors demoted may be less than what is specified by the host.

Table 91 - HYBRID DEMOTE BY SIZE - Number of logical sectors affected

Sector Count[a] Is less than Actual[b]	**Number of Logical Sectors to change to the To Priority[c]**
Y	Sector Count[a]
N	Actual[b] number

[a] The value of the Sector Count field in the command parameters.

[b] The actual number of logical sectors in the non-volatile caching medium associated with the From priority.

[c] The value of the Hybrid Priority field (see 13.20.6.5) within the Hybrid Information field (see 13.20.6).

The device selects which logical sectors are to be changed based on vendor specific criteria (e.g., age and usage).

The device shall return command aborted if:

a) the value of the FromPriority field is less than or equal to the value of the Hybrid Priority field; or

b) the value of the FromPriority field is equal to the Maximum Hybrid Priority Level field (see 13.7.5.4.8) in the Hybrid Information log and the Max Priority Behavior bit (see 13.7.5.4.11) is set to one.

The device shall complete the operation before returning command complete.

Field	**7**	**6**	**5**	**4**	**3**	**2**	**1**	**0**
Features(7:0)	FromPriority				2h			

Features(15:8)	Sector Count (7:0)				
Count(7:0)	TAG				Reserved
Count(15:8)	Sector Count (15:8)				
LBA(7:0)	Sector Count (23:16)				
LBA(15:8)	Sector Count (31:24)				
LBA(23:16)	Reserved				
LBA(31:24)	Reserved				
LBA(39:32)	Reserved				
LBA(47:40)	Reserved				
ICC	Reserved				
Auxiliary(7:0)	Reserved				
Auxiliary(15:8)	Reserved				
Auxiliary(23:16)	Hybrid Information				
Auxiliary(31:24)	Reserved				
Device	Res	1	Res	0	Reserved
Command	**63h**				

Figure 241+1 –HYBRID DEMOTE BY SIZE - Command definition

FromPriority The FromPriority specifies the Hybrid Priority level corresponding to the data that is to be changed to the hybrid priority specified in the Hybrid Information field.

Sector Count The Sector Count field specifies the requested number of logical sectors to which the change of hybrid priority applies.

TAG The TAG value shall be assigned by host software to be different from all other TAG values corresponding to outstanding commands. TAG shall not exceed the value specified in IDENTIFY DEVICE word 75.

Hybrid Information The Hybrid Information field contains the data hints for the specified logical sectors (see 13.20.6).

13.6.5.4.8 [Editor's note 13.6.6.4.2]Success Outputs

If a HYBRID DEMOTE BY SIZE command completes without error, then a Set Device Bits FIS shall be sent to the host. This SDB FIS may also indicate other completed commands.

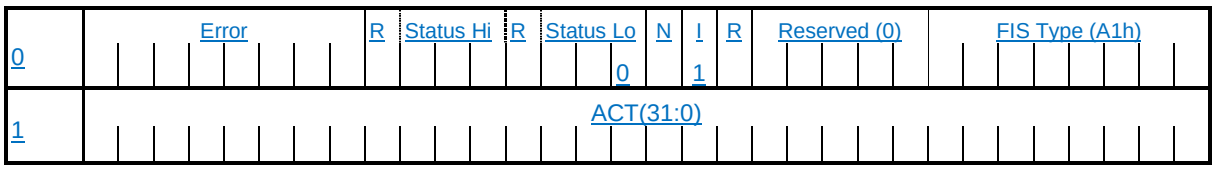

Figure 241+2 –HYBRID DEMOTE BY SIZE - Successful completion

ACT The ACT field of the Set Device Bits FIS communicates completion notification for each of up to 32 commands. The field is bit-significant and the device sets bit positions to one for each command tag it is indicating completion notification for.

The device may set more than one bit to one if it is explicitly aggregating successful status returns. The device shall set to one the bit associated with the TAG value for the HYBRID DEMOTE BY SIZE command.

Error The Error register shall contain 00h.

Status As defined in [Editor's note 10.3.6]. The ERR bit shall be cleared to zero indicating successful command completion. Bit 4 may be set to one.

I Interrupt bit. The interrupt bit shall be set to one.

All other fields as defined in [Editor's note 10.3.6].

13.6.5.4.9 [Editor's note 13.6.6.4.3]Error Outputs

13.6.5.5.9.2 [Editor's note 13.6.6.4.3.1]Upon receipt of a command

If the device has received a command that has not yet been acknowledged by clearing the BSY bit to zero and an error is encountered, the device shall transmit a Register Device to Host FIS (see [Editor's note Figure 241+3]) with the ERR bit set to one and the BSY bit cleared to zero in the Status field, the ATA error code in the Error field.

Field	**7**	**6**	**5**	**4**	**3**	**2**	**1**	**0**
Error	ERROR							
Count(7:0)	na							
Count(15:8)	na							
LBA(7:0)	na							
LBA(15:8)	na							
LBA(23:16)	na							
LBA(31:24)	na							
LBA(39:32)	na							
LBA(47:40)	na							
Device	na							
Status	BSY	DRDY	DF	na	DRQ	na	na	ERR

Figure 241+3 –HYBRID DEMOTE BY SIZE - error on command receipt

ERROR ATA error code for the failure condition of the failed command

BSY 0

DRDY 1

DF 0

DRQ 0

ERR 1

Status bit 4 may be set to one.

Following transmission of the Register Device to Host FIS, the device shall stop processing any outstanding or new commands until the Queued Error Log (see [Editor's note 13.7.7]) has been read before continuing to abort all outstanding commands. See [Editor's note 13.6.3.3] for more details.

13.6.5.5.9.2 [Editor's note 13.6.6.4.3.2]During execution of a command

If all commands have been acknowledged by clearing the BSY bit to zero and an error condition is detected, the device shall transmit a Set Device Bits FIS (see Figure 241+4) to the host with the ERR bit set to one in the Status field, the ATA error code in the Error field, and the Interrupt bit set to one. All outstanding commands at the time of an error are aborted as part of the error response and may be re-issued as appropriate by the host. For any commands that have not completed or have completed with error, the device shall clear the corresponding ACT bits to zero in the Set Device Bits FIS.

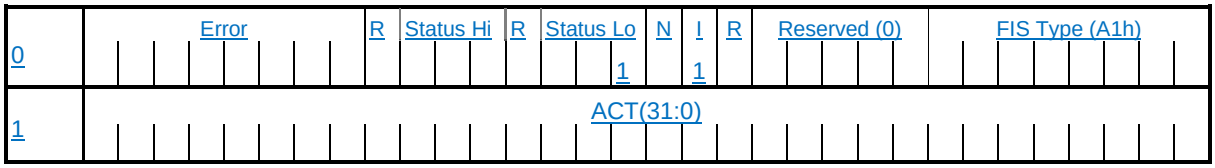

Figure 241+4 –HYBRID DEMOTE BY SIZE – error during execution

ACT — The ACT field of the Set Device Bits FIS communicates successful completion notification for each of up to 32 queued commands. The field is bit-significant and the device sets bit positions to one for each command tag it is indicating successful completion notification for. The device may set more than one bit to one if it is explicitly aggregating successful status returns.

Error — The Error register shall contain the ATA error code.

Status — As defined in [editors note: 10.3.6xxx]. The ERR bit shall be set to one indicating an NCQ error has occurred. Status bit 4 may be set to one.

I — Interrupt bit. The interrupt bit shall be set to one.

All other fields as defined in [Editor's note 10.3.6].

Only the registers that are updated as part of the Set Device Bits FIS are modified if the device signals an error condition when the BSY bit in the shadow Status register is cleared to zero, leaving the other Shadow Register Block Registers unchanged. If the device signals an error condition when the BSY bit in the shadow Status register is set to one, the device clears the BSY bit to zero with a Register Device to Host FIS which updates all registers in the Shadow Register Block.

Following transmission of the Set Device Bits FIS, the device shall stop processing any outstanding or new commands until the Queued Error Log (see [Editor's note 13.7.7]) has been read before continuing to abort all outstanding commands. See [Editor's note 13.6.3.3] for more details.

[editors note: add the HYBRID CHANGE BY LBA RANGE subcommand to the NCQ NON-DATA command]

13.6.5.5 [Editor's note 13.6.6.5]HYBRID CHANGE BY LBA RANGE Subcommand

13.6.5.4.10 [Editor's note 13.6.6.5.1]HYBRID CHANGE BY LBA RANGE Subcommand definition

The HYBRID CHANGE BY LBA RANGE command is used to change the hybrid information associated with a specified range of logical sectors. Support for this subcommand is indicated in the NCQ NON-DATA Log (see 13.7.9.14).

The device shall set the hybrid priority associated with a number of logical sectors starting at the LBA specified by the Starting LBA field, regardless of what hybrid priority is associated with the selected logical sectors. The requested new hybrid priority may be any valid hybrid priority.

The Sector Count specifies the number of logical sectors for which the device should change the hybrid priority to the value specified in the Hybrid Priority field in the Hybrid Information field.

If the Max Priority Behavior bit (see 13.7.5.4.11) is set to 1 and the Hybrid Priority field is equal to the maximum hybrid priority, then the device shall copy all of the selected logical sectors that are not already in the non-volatile caching medium into the non-volatile caching medium. Otherwise, if the Hybrid priority field is non-zero, then the device should copy any of the selected logical sectors that are not already in the non-volatile caching medium into the non-volatile caching medium.

If:

a) the Hybrid Priority field (see 13.20.6.5) is set to the Maximum Priority;

b) the MaxPriority Behavior bit (see 13.7.5.4.11) is set to one; and

c) the non-volatile caching medium does not have mapping resources,

then:

a) the device shall return command aborted; and

b) in the Queued Error log, the device shall:

   A. set the the Sense Key field to ABORTED COMMAND; and

   B. set the the additional sense code (i.e., ASC field and ASCQ field) to INSUFFICIENT RESOURCES.

If any of the selected logical sectors are already in the non-volatile caching medium associated with other Hybrid Information field values, then the new values shall replace the previous values.

The device shall complete the requested changes and move the data, if needed, before returning command completed.

NOTE: The device may take 30 seconds or more to complete the command if a large LBA range is specified. To minimize system response issues, large LBA ranges should be broken up into multiple smaller operations.

If the value of the Hybrid Priority field is zero, then the device may evict the selected logical sectors from the non-volatile caching medium.

Field	7	6	5	4	3	2	1	0
Features(7:0)	Reserved				3h			

Features(15:8)	Sector Count (7:0)				
Count(7:0)	TAG				Reserved
Count(15:8)	Sector Count (15:8)				
LBA(7:0) LBA(15:8) LBA(23:16) LBA(31:24) LBA(39:32) LBA(47:40)	StartingLBA				
ICC	Reserved				
Auxiliary(7:0)	Reserved				
Auxiliary(15:8)	Reserved				
Auxiliary(23:16)	Hybrid Information				
Auxiliary(31:24)	Reserved				
Device	Res	1	Res	0	Reserved
Command	**63h**				

Figure 241+5 –HYBRID CHANGE BY LBA RANGE - Command definition

Starting LBA The Starting LBA field specifies the first LBA.

Sector Count The Sector Count field specifies the requested number of logical sectors, starting from Starting LBA.

Hybrid Information The Hybrid Information field contains the data hints for the specified logical sectors (see 13.20.6).

TAG The TAG value shall be assigned by host software to be different from all other TAG values corresponding to outstanding commands. TAG shall not exceed the value specified in IDENTIFY DEVICE word 75.

13.6.5.4.11 [Editor's note 13.6.6.5.2]Success Outputs

If a HYBRID CHANGE BY LBA RANGE command completes without error, then a Set Device Bits FIS shall be sent to the host. This SDB FIS may also indicate other completed commands.

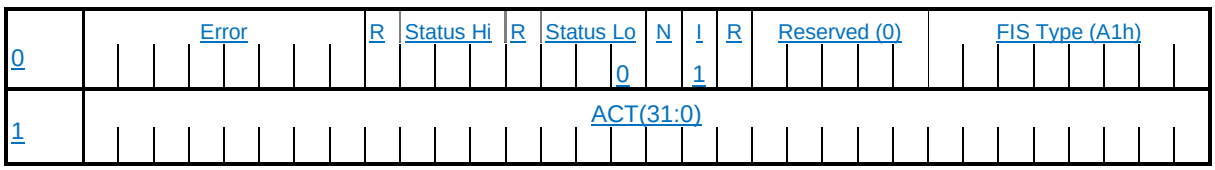

Figure 241+6 –HYBRID CHANGE BY LBA RANGE - Successful completion

ACT The ACT field of the Set Device Bits FIS communicates completion notification for each of up to 32 commands. The field is bit-significant and the device sets bit positions to one for each command tag it is indicating completion notification for.

The device may set more than one bit to one if it is explicitly aggregating successful status returns. The device shall set to one the bit associated with the TAG value for the HYBRID CHANGE BY LBA RANGE command.

Error — The Error register shall contain 00h.

Status — As defined in [Editor's note 10.3.6]. The ERR bit shall be cleared to zero indicating successful command completion. Bit 4 may be set to one.

I — Interrupt bit. The interrupt bit shall be set to one.

All other fields as defined in [Editor's note 10.3.6].

13.6.5.4.12 [Editor's note 13.6.6.5.3]Error Outputs

13.6.5.5.12.2 [Editor's note 13.6.6.5.3.1]Upon receipt of a command

If the device has received a command that has not yet been acknowledged by clearing the BSY bit to zero and an error is encountered, the device shall transmit a Register Device to Host FIS (see [Editor's note Figure 214]) with the ERR bit set to one and the BSY bit cleared to zero in the Status field, the ATA error code in the Error field.

[editors note: the LBA fields have been re-sorted in bit order]

Field	7	6	5	4	3	2	1	0
Error	ERROR							
Count(7:0)	na							
Count(15:8)	na							
LBA(7:0)	na							
LBA(15:8)	na							
LBA(23:16)	na							
LBA(31:24)	na							
LBA(39:32)	na							
LBA(47:40)	na							
Device	na							
Status	BSY	DRDY	DF	na	DRQ	na	na	ERR

Figure 241+7 –HYBRID CHANGE BY LBA RANGE - error on command receipt

ERROR — ATA error code for the failure condition of the failed command

BSY — 0

DRDY — 1

DF — 0

DRQ — 0

ERR — 1

Status bit 4 may be set to one.

Following transmission of the Register Device to Host FIS, the device shall stop processing any outstanding or new commands until the Queued Error Log (see [Editor's note 13.7.7]) has been read before continuing to abort all outstanding commands. See [Editor's note 13.6.3.3] for more details.

13.6.5.5.12.2 [Editor's note 13.6.6.4.3.2]During execution of a command

If all commands have been acknowledged by clearing the BSY bit to zero and an error condition is detected, the device shall transmit a Set Device Bits FIS (see [Editor's note Figure 241+8]) to the host with the ERR bit set to one in the Status field, the ATA error code in the Error field, and the Interrupt bit set to one. All outstanding commands at the time of an error are aborted as part of the error response and may be re-issued as appropriate by the host. For any commands that have not completed or have completed with error, the device shall clear the corresponding ACT bits to zero in the Set Device Bits FIS.

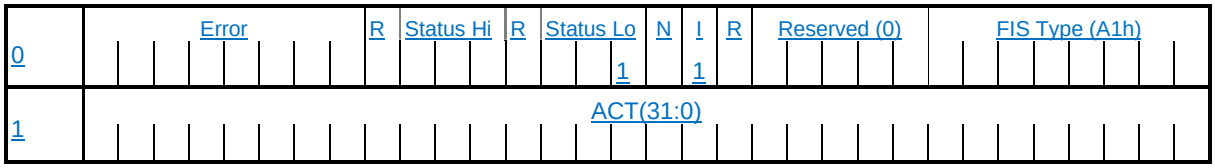

Figure 241+8 –HYBRID CHANGE BY LBA RANGE – error during execution

ACT — The ACT field of the Set Device Bits FIS communicates successful completion notification for each of up to 32 queued commands. The field is bit-significant and the device sets bit positions to one for each command tag it is indicating successful completion notification for. The device may set more than one bit to one if it is explicitly aggregating successful status returns.

Error — The Error register shall contain the ATA error code.

Status — As defined in [editors note: 10.3.6xxx]. The ERR bit shall be set to one indicating an NCQ error has occurred. Status bit 4 may be set to one.

I — Interrupt bit. The interrupt bit shall be set to one.

All other fields as defined in [Editor's note 10.3.6].

Only the registers that are updated as part of the Set Device Bits FIS are modified if the device signals an error condition when the BSY bit in the shadow Status register is cleared to zero, leaving the other Shadow Register Block Registers unchanged. If the device signals an error condition when the BSY bit in the shadow Status register is set to one, the device clears the BSY bit to zero with a Register Device to Host FIS which updates all registers in the Shadow Register Block.

Following transmission of the Set Device Bits FIS, the device shall stop processing any outstanding or new commands until the Queued Error Log (see [Editor's note 13.7.7]) has been read before continuing to abort all outstanding commands. See [Editor's note 13.6.3.3] for more details.

[editors note: add the HYBRID CONTROL subcommand to the NCQ NON-DATA command]

13.6.5.6 [Editor's note 13.6.6.6]HYBRID CONTROL Subcommand

13.6.5.4.13 [Editor's note 13.6.6.6.1]HYBRID CONTROL Subcommand definition

The HYBRID CONTROL command provides parameters for the use of the non-volatile caching medium. Support for this subcommand is indicated in the NCQ NON-DATA Log (see 13.7.9.15).

In contrast, the SET FEATURES Enable/Disable Hybrid Information subcommand (see 13.3.15) provides a mechanism to enable or disable the Hybrid Information feature.

If the command completes without error, then the device shall preserve the values of the Dirty Low Threshold field and the Dirty High Threshold field across all resets and power cycle events. Current values are available in the Hybrid Information log (see 13.7.12).

Field	**7**	**6**	**5**	**4**	**3**	**2**	**1**	**0**
Features(7:0)	Disable Caching Medium	Reserved			4h			
Features(15:8)	Reserved							
Count(7:0)	TAG					Reserved		
Count(15:8)	Reserved							
LBA(7:0)	Dirty Low Threshold							
LBA(15:8)	Dirty High Threshold							
LBA(23:16)	Reserved							
LBA(31:24)								
LBA(39:32)								
LBA(47:40)								
ICC	Reserved							
Auxiliary(7:0)	Reserved							
Auxiliary(15:8)	Reserved							
Auxiliary(23:16)	Reserved							
Auxiliary(31:24)	Reserved							
Device	Res	1	Res	0	Reserved			
Command	**63h**							

Figure 241+9 –HYBRID CONTROL- Command definition

Disable Caching Medium

If the Disable Caching Medium field is cleared to zero, then the device shall process the Dirty Low Threhold field and the Dirty High Threshold field.

If the Disable Caching Medium field is set to one and the Hybrid Information feature is enabled, then the device shall:

1. ignore the contents of the Dirty Low Threhold field and the Dirty High Threshold field;
2. change the value of the Enabled field (see 13.7.5.4.3) of the Hybrid Information log to 80h (i.e., Hybrid Information Disable In Process);
3. sync all dirty data in the non-volatile caching medium to the primary medium;
4. evict all data from the non-volatile caching medium;
5. disable the Hybrid Information feature (see 13.3.15);
6. change the value of the Enabled field (see 13.7.5.4.3) of the Hybrid Information log to 00h (i.e., Hybrid Information Disabled); and
7. disable the use of the non-volatile caching medium for storing user data until the device processes a SET FEATURES Enable Hybrid Information subcommand.

If the device processes a reset or a power cycle while the value of the Enabled field (see 13.7.5.4.3) of the Hybrid Information log is 80h (i.e., Hybrid Information Disable In Process), then the device shall change the value of the Enabled field to FFh (i.e., Hybrid Information Enabled).

If the Disable Caching Medium bit is set to one, then the device may report command completion before making the requested changes. The host should check the Hybrid Information log to determine if the requested changes have been completed.

Dirty Low Threshold

The Dirty Low Threshold represents the threshold for the amount of dirty user logical sectors in the non-volatile caching medium at which point syncing operations should stop. The value of the Dirty Low Threshold field divided by 255 specifies a fraction of the total reported NVM Size of the non-volatile caching medium that contains dirty logical sectors. The device shall preserve this setting over all power cycles and all resets. See 13.20.7xxx for additional information on syncing. The current value is reported in the Hybrid Information log (see 13.7.5.4.5xxx).

Dirty High Threshold

The Dirty High Threshold represents the threshold for the amount of dirty user logical sectors in the non-volatile caching medium at which point syncing operations should begin. The value of the Dirty High Threshold field divided by 255 specifies a fraction of the total reported NVM Size of the non-volatile caching medium that contains dirty data. The device shall preserve this setting over all power cycles and all resets. The current value is reported in the Hybrid Information log (see 13.7.5.4.6xxx). See 13.20.7xxx for additional information on syncing.

TAG

The TAG value shall be assigned by host software to be different from all other TAG values corresponding to outstanding commands. TAG shall not exceed the value specified in IDENTIFY DEVICE word 75.

13.6.5.4.14 [Editor's note 13.6.6.6.2]Success Outputs

If a HYBRID CONTROL command completes without error, a Set Device Bits FIS shall be sent to the host. This SDB FIS may also indicate other completed commands.

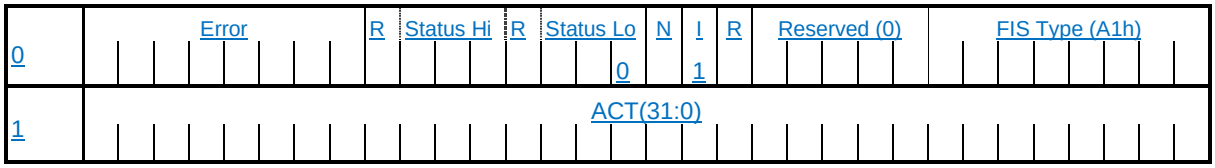

Figure 241+10 –HYBRID CONTROL - Successful completion

ACT — The ACT field of the Set Device Bits FIS communicates completion notification for each of up to 32 commands. The field is bit-significant and the device sets bit positions to one for each command tag it is indicating completion notification for. The device may set more than one bit to one if it is explicitly aggregating successful status returns. The device shall set to one the bit associated with the TAG value for the HYBRID CONTROL command.

Error — The Error register shall contain 00h.

Status — As defined in [editors note: 10.3.6xxx]. The ERR bit shall be cleared to zero indicating successful command completion. Bit 4 may be set to one.

I — Interrupt bit. The interrupt bit shall be set to one.

All other fields as defined in [Editor's note 10.3.6].

13.6.5.4.15 [Editor's note 13.6.6.6.3]Error Outputs

13.6.5.6.3.1 [Editor's note 13.6.6.6.3.1]Upon receipt of a command

If the device has received a command that has not yet been acknowledged by clearing the BSY bit to zero and an error is encountered, the device shall transmit a Register Device to Host FIS (see [Editor's note Figure 241+11]) with the ERR bit set to one and the BSY bit cleared to zero in the Status field, the ATA error code in the Error field.

[editors note: the LBA fields have been re-sorted in bit order]

Field	7	6	5	4	3	2	1	0
Error	ERROR							
Count(7:0)	na							
Count(15:8)	na							
LBA(7:0)	na							
LBA(15:8)	na							
LBA(23:16)	na							
LBA(31:24)	na							
LBA(39:32)	na							
LBA(47:40)	na							
Device	na							
Status	BSY	DRDY	DF	na	DRQ	na	na	ERR

Figure 241+11 –HYBRID CONTROL - error on command receipt

ERROR ATA error code for the failure condition of the failed command

BSY 0

DRDY 1

DF 0

DRQ 0

ERR 1

Status bit 4 may be set to one.

Following transmission of the Register Device to Host FIS, the device shall stop processing any outstanding or new commands until the Queued Error Log (see [Editor's note 13.7.7]) has been read before continuing to abort all outstanding commands. See [Editor's note 13.6.3.3] for more details.

13.6.5.6.3.2 [Editor's note 13.6.6.6.3.2]During execution of a command

If all commands have been acknowledged by clearing the BSY bit to zero and an error condition is detected, the device shall transmit a Set Device Bits FIS (see [Editor's note Figure 241+12) to the host with the ERR bit set to one in the Status field, the ATA error code in the Error field, and the Interrupt bit set to one. All outstanding commands at the time of an error are aborted as part of the error response and may be re-issued as appropriate by the host. For any commands that have not completed or have completed with error, the device shall clear the corresponding ACT bits to zero in the Set Device Bits FIS.

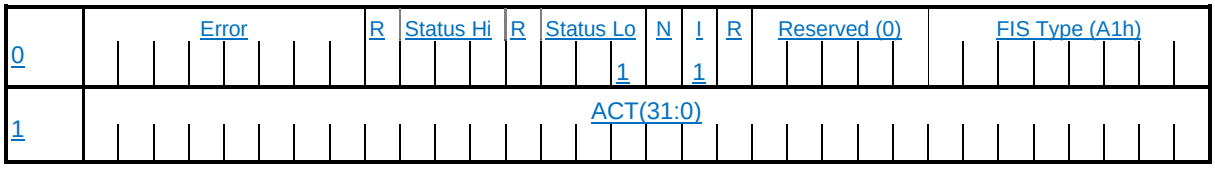

Figure 241+12 –HYBRID CONTROL– error during execution

ACT	The ACT field of the Set Device Bits FIS communicates successful completion notification for each of up to 32 queued commands. The field is bit-significant and the device sets bit positions to one for each command tag it is indicating successful completion notification for. The device may set more than one bit to one if it is explicitly aggregating successful status returns.
Error	The Error register shall contain the ATA error code.
Status	As defined in [editors note: 10.3.6xxx]. The ERR bit shall be set to one indicating an NCQ error has occurred. Status bit 4 may be set to one.
I	Interrupt bit. The interrupt bit shall be set to one.

All other fields as defined in [Editor's note 10.3.6].

Only the registers that are updated as part of the Set Device Bits FIS are modified if the device signals an error condition when the BSY bit in the shadow Status register is cleared to zero, leaving the other Shadow Register Block Registers unchanged. If the device signals an error condition when the BSY bit in the shadow Status register is set to one, the device clears the BSY bit to zero with a Register Device to Host FIS which updates all registers in the Shadow Register Block.

Following transmission of the Set Device Bits FIS, the device shall stop processing any outstanding or new commands until the Queued Error Log (see [Editor's note 13.7.7]) has been read before continuing to abort all outstanding commands. See [Editor's note 13.6.3.3] for more details.

13.6.12 ~~<13.6.6>~~[Editor's note 13.6.7]RECEIVE FPDMA QUEUED

[editors note: make no changes]

13.6.13 ~~<13.6.7>~~[Editor's note 13.6.8]SEND FPDMA QUEUED

13.6.13.4 [Editor's note 13.6.8.1]SEND FPDMA QUEUED definition

The 512 Byte Block DMA OUT subcommands make use of this transport command. The SEND FPDMA QUEUED command supports LBA mode only and uses 48-bit addressing only. The format of the command is defined in ~~Figure 257~~Figure 246.

[editors note: the LBA fields have been re-sorted in bit order]

13.6.13.5 ~~<13.6.7.1>~~[Editor's note 13.6.8.2]Inputs

~~Register~~ Field	7	6	5	4	3	2	1	0
Features (7:0)	Sector Count (7:0)							
Features (15:8)	Sector Count (15:8)							
Count (7:0)	TAG					Reserved		
Count (15:8)	PRIO 1:0		Res	Subcommand				
LBA ~~Low~~ (7:0)	~~LBA (7:0)~~ Subcommand Specific							
LBA ~~Mid~~ (15:8)	~~LBA (15:8)~~ Subcommand Specific							
LBA ~~High~~ (23:16)	~~LBA (23:16)~~ Subcommand Specific							
LBA ~~Low~~ (31:24)	~~LBA (31:24)~~ Subcommand Specific							
LBA ~~Mid~~ (39:32)	~~LBA (39:32)~~ Subcommand Specific							
LBA ~~High~~ (47:40)	~~LBA (47:40)~~ Subcommand Specific							
ICC(7:0)	ICC (7:0)							
Auxiliary(7:0)	~~Auxiliary (7:0)~~Subcommand Specific							
Auxiliary (15:8)	~~Auxiliary (15:8)~~Subcommand Specific							
Auxiliary(23:16)	Subcommand Specific							
Auxiliary(31:24)	~~Reserved~~Subcommand Specific							
Device	Res	1	Res	0	Reserved			
Command	**64h**							

Figure 246 – SEND FPDMA QUEUED command definition

Sector Count The number of 512 byte blocks to be transferred, 0000h indicates that 65, 536 512-byte blocks are to be transferred.

TAG The TAG value shall be assigned by host software to be different from all other TAG values corresponding to outstanding commands. The assigned TAG value shall not exceed the value specified in IDENTIFY DEVICE word 75.

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 and Isochronous priority requests. The device shall make a best effort to complete Isochronous priority requests prior to its associated deadline. The Priority values are defined as follows:

a) 00b Normal Priority;

b) 01b Isochronous – deadline dependent priority;

c) 10b High priority; and

d) 11b Reserved.

Subcommand — Subcommand specific (see 13.6.13.8)

LBA — Subcommand specific (see 13.6.13.8)

ICC — The Isochronous Command Completion (ICC) field shall be assigned by the host based on the intended deadline associated with the command issued. By default, ~~when~~if deadline is expired, the device shall continue to complete the command as soon as possible.

Auxiliary — Subcommand specific (see 13.6.13.8)

Upon accepting the command, the device shall clear the BSY bit by transmitting a Register Device to Host FIS to the host with the BSY bit cleared to zero in the Status field of the FIS. The ability for the device to quickly clear the BSY bit allows the host to issue another queued command without blocking on this bit.

The host shall check the BSY bit in the shadow Status register before attempting to issue a new command in order to determine that the device is ready to receive another command (and determine that the host has write access to the Shadow Register Block Registers).

The device shall not trigger an interrupt in response to having successfully received the command, so the initial status return that clears BSY shall not have an interrupt associated with it.

13.6.13.6 ~~<13.6.7.2>~~[Editor's note 13.6.8.2]Success Outputs

Upon successful completion of one or more outstanding commands, the device shall transmit a Set Device Bits FIS with the Interrupt bit set to one and one or more bits set to one in the ACT field corresponding to the bit position for each command TAG that has completed since the last status notification was transmitted. The ERR bit in the Status register shall be cleared to zero and the value in the Error register shall be zero.

The ACT field occupies the last 32 bits of the Set Device Bits FIS as defined in Figure 247.

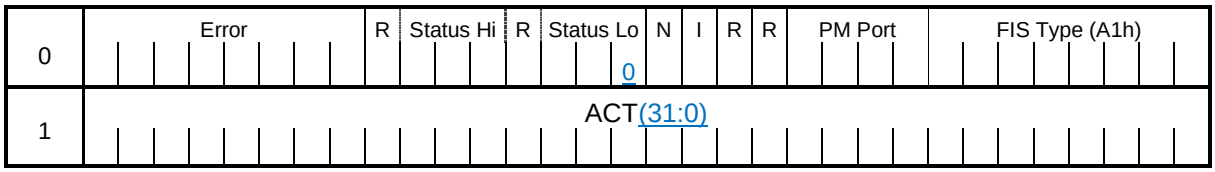

Figure 247 – Set Device Bits FIS for successful SEND FPDMA QUEUED command completion

ACT — The ACT field of the Set Device Bits FIS communicates successful completion notification for each of up to 32 queued commands. The field is bit-significant and the device sets bit positions to one for each command tag it is indicating

successful completion notification for. The device may set more than one bit to one if it is explicitly aggregating successful status returns.

Error — The Error register shall be cleared to zero.

Status — As defined in ~~section~~ [Editor's note 10.3.6]. The ERR bit shall be cleared to zero indicating successful command completion. Bit 4 may be set to one.

I — Interrupt bit. The interrupt bit shall be set to one.

All other fields as defined in ~~section~~ [Editor's note 10.3.6].

Devices should be aware that if choosing to aggregate status to the point where many of the outstanding commands have actually completed successfully without notification to the host, that an error may cause the final completion status of those commands to be failure. The device should be selective ~~when~~if using status aggregation for outstanding queued commands to ensure the host is made aware of successful completion for outstanding commands in a way that an error ~~would not~~be unable to force a high number of unnecessary command retries.

13.6.13.7 ~~<13.6.7.3>~~[Editor's note 13.6.8.3]Error Outputs

13.6.5.4.1 [Editor's note 13.6.8.3.1]Error Outputs status

If the device has received a command that has not yet been acknowledged by clearing the BSY bit to zero and an error is encountered, the device shall transmit a Register Device to Host FIS (see Figure 248) with the ERR bit set to one and the BSY bit cleared to zero in the Status field, the ATA error code in the Error field.

[editors note: the LBA fields have been re-sorted in bit order]

~~Register~~ Field	7	6	5	4	3	2	1	0
Error	ERROR							
Count (7:0)	na							
Count (15:8)	na							
LBA ~~Low~~ (7:0)	na							
LBA ~~Mid~~ (15:8)	na							
LBA ~~High~~ (23:16)	na							
LBA ~~Low~~ (31:24)	na							
LBA ~~Mid~~ (39:32)	na							
LBA ~~High~~ (47:40)	na							
Device	na							
Status	BSY	DRDY	DF	na	DRQ	na	na	ERR

Figure 248 – SEND FPDMA QUEUED error status result values on command receipt

ERROR ATA error code for the failure condition of the failed command

BSY 0

DRDY 1

DF 0

DRQ 0

ERR 1

Status bit 4 may be set to one.

Following transmission of the Register Device to Host FIS, the device shall stop processing any outstanding or new commands until the Queued Error Log (see [Editor's note 13.7.7]) has been read before continuing to abort all outstanding commands. See [Editor's note 13.6.3.3] for more details.

13.6.5.4.2 ~~<13.6.7.3.1>~~[Editor's note 13.6.8.3.2]During Execution of a Command

If all commands have been acknowledged by clearing the BSY bit to zero and an error condition is detected, the device shall transmit a Set Device Bits FIS (see [Editor's note Figure 260]) to the host with the ERR bit set to one in the Status field, the ATA error code in the Error field, and the Interrupt bit set to one. All outstanding commands at the time of an error shall be aborted as part of the error response and may be re-issued as appropriate by the host. For any commands that have not completed successfully or have resulted in error, the device shall clear the corresponding ACT bits to zero in the Set Device Bits FIS.

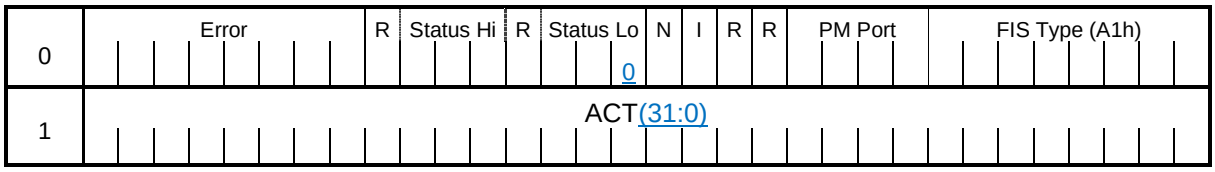

Figure 249 – Set Device Bits FIS with error notification, and command completions

ACT The ACT field of the Set Device Bits FIS communicates successful completion notification for each of up to 32 queued commands. The field is bit-significant and the device sets bit positions to one for each command tag it is indicating successful completion notification for. The device may set more than one bit to one if it is explicitly aggregating successful status returns.

Error The Error register shall contain the ATA error code.

Status As defined in [Editor's note 10.3.6]. The ERR bit shall be set to one indicating an NCQ error has occurred. Status bit 4 may be set to one.

I Interrupt bit. The interrupt bit shall be set to one.

All other fields as defined in section [Editor's note 10.3.6].

Only the registers that are updated as part of the Set Device Bits FIS are modified if the device signals an error condition ~~when~~if the BSY bit in the shadow Status register is cleared to zero, leaving the other Shadow Register Block Registers unchanged. If the device signals an error condition ~~when~~if the BSY bit in the shadow Status register is set to one, the device clears the BSY bit to zero with a Register Device to Host FIS which updates all registers in the Shadow Register Block.

Following transmission of the Set Device Bits FIS, the device shall stop processing any outstanding or new commands until the Queued Error Log (see [Editor's note 13.7.7]) has been before continuing to abort all outstanding commands. See [Editor's note 13.6.3.3] for more details.

13.6.13.8 ~~<13.6.7.4>~~[Editor's note 13.6.8.4]SEND FPDMA QUEUED Subcommands

Subcommands for the SEND FPDMA QUEUED commands are defined in Table 93 and are contained within the Count (13:8) field. Support for each subcommand is reported in the NCQ Send and Receive log (see 13.7.10).

Table 92 – Subcommands for SEND FPDMA QUEUED

Value	Subcommand
00h	Data Set Management (see ~~section~~ 13.6.13.9)
01h	Hybrid Evict (see 13.6.13.10)
02h ~~01h~~ ..1Fh	Reserved

13.6.13.9 ~~<13.6.7.5>~~[Editor's note 13.6.8.5]DATA SET MANAGEMENT

13.6.5.4.1 [Editor's note 13.6.8.5.1]DATA SET MANAGEMENT Subcommand

The DATA SET MANAGEMENT subcommand functionality and behavior is dependent on all requirements of the DATA SET MANAGEMENT command and the IDENTIFY DEVICE command defined in ACS-2.

[editors note: the LBA fields have been re-sorted in order of bit position]

13.6.5.4.2 ~~<13.6.7.5.1>~~[Editor's note 13.6.8.5.2]~~Inputs~~Subcommand specific parameters for the DATA SET MANAGEMENT command

~~Register~~ Field	7	6	5	4	3	2	1	0
LBA~~-~~(7:0)	Reserved							
LBA~~-~~(15:8)	Reserved							
LBA~~-~~(23:16)	Reserved							
LBA~~-~~(31:24)	Reserved							
LBA~~-~~(39:32)	Reserved							
LBA~~-~~(47:40)	Reserved							
Auxiliary~~-~~(7:0)	Reserved							Trim
Auxiliary~~-~~(15:8)	Reserved							
Auxiliary~~-~~(23:16)	Reserved							
Auxiliary~~-~~(31:24)	Reserved							
~~LBA (7:0)~~	~~Reserved~~							
~~LBA (15:8)~~	~~Reserved~~							
~~LBA (23:16)~~	~~Reserved~~							
~~LBA (31:24)~~	~~Reserved~~							
~~LBA (39:32)~~	~~Reserved~~							
~~LBA (47:40)~~	~~Reserved~~							

Figure 250 – ~~SEND FPDMA QUEUED, Subcommand = 00h~~*Subcommand specific parameters for the DATA SET MANAGEMENT command*

Trim As defined by the DATA SET MANAGEMENT command in ACS-2.

*13.6.5.4.3 ~~<13.6.7.5.2>~~[Editor's note *13.6.8.5.3*]Success Outputs*

See 13.6.13.6

*13.6.5.4.4 ~~<13.6.7.5.3>~~[Editor's note *13.6.8.5.4*]Error Outputs*

See 13.6.13.7

*13.6.5.4.5 ~~<13.6.7.5.4>~~[Editor's note *13.6.8.5.5*]Output from the Host to the Device Data Structure*

As defined in the DATA SET MANAGEMENT command in ACS-2.

[editors note: add the HYBRID EVICT subcommand to the SEND FPDMA QUEUED command]

13.6.13.10 [Editor's note 13.6.8.6]HYBRID EVICT Subcommand

13.6.5.4.1 [Editor's note 13.6.8.6.1] HYBRID EVICT Subcommand definition

The HYBRID EVICT subcommand evicts data from the non-volatile caching medium to the primary medium. Support for this subcommand is indicated in the NCQ Send And Receive Log (see 13.7.10.5).

If the Evict All field (see [Editor's note Figure 250+1]) is cleared to zero, then for each LBA range specified in the data transferred from the host (see 13.6.5.4.4), the device:

a) shall sync all dirty data in the non-volatile caching medium;

b) may evict the specified LBA range in the non-volatile caching medium; and

c) may free mapping resources in the non-volatile caching medium related to the specified LBA range.

If the Evict All field is set to one, then:

a) the device shall ignore all data transferred from the host (see 13.6.5.4.4); and

b) for all data in the non-volatile caching medium, the device shall:

   A) sync all dirty data in the non-volatile caching medium;

   B) evict all user data in the non-volatile caching medium; and

   C) free mapping resources in the non-volatile caching medium.

If the value of the Sector Count field (see Figure 246) is greater than the value of the Maximum Eviction Data Blocks field (see 13.7.5.4.16) of the Hybrid Information log, then the device shall return command aborted.

The device may limit the number of HYBRID EVICT commands that are in the queue at the same time. The device shall return command aborted if:

a) a new HYBRID EVICT command is accepted; and

b) the number of HYBRID EVICT commands in the queue was previously equal to the value of the Maximum Eviction Commands field (see 13.7.5.4.15) of the Hybrid Information log.

If the device processes any reset while processing a HYBRID EVICT command, then the resulting condition of the non-volatile caching medium is indeterminate.

The device should return command completion with no error (see 13.6.5.4.2) after all selected LBA ranges have been successfully evicted.

NOTE: This command may take longer to complete than a typical maximum timeout.

See the SEND FPDMA QUEUED command (see 13.6.13) for the contents of the Count field.

13.6.5.4.2 Subcommand specific parameters for the HYBRID EVICT command

Field	7	6	5	4	3	2	1	0
LBA (7:0)	Reserved							
LBA (15:8)	Reserved							
LBA (23:16)	Reserved							
LBA (31:24)	Reserved							
LBA (39:32)	Reserved							
LBA (47:40)	Reserved							
Auxiliary (7:0)	Reserved							Evict All
Auxiliary (15:8)	Reserved							
Auxiliary (23:16)	Reserved							
Auxiliary (31:24)	Reserved							

Figure 250+1 - ~~SEND FPDMA QUEUED, Subcommand = 01h~~*Subcommand specific parameters for the HYBRID EVICT command*

Evict All The Evict All field specifies that all of the data in the non-volatile caching medium is to be evicted

13.6.5.4.3 [Editor's note 13.6.8.6.2]*Success Outputs*

See 13.6.13.6

13.6.5.4.4 [Editor's note 13.6.8.6.3]Error Outputs

See 13.6.13.7

13.6.5.4.5 [Editor's note 13.6.8.6.4]Output from the Host to the Device Data Structure

[Editor's note Figure 250+2] describes the format for all 512-byte data blocks transferred from the host to the device for the HYBRID EVICT command, containing up to 64 LBA Range entries (see [Editor's note Figure 250+3]). There may be more than one 512-byte data block transferred. The LBA Range entries shall be sorted in order of increasing Starting LBA. If the value of the Range Length field of an LBA Range entry is cleared to zero, then the device shall ignore the LBA Range entry and all following LBA Range entries.

For any LBA range, if the Starting LBA plus the Range Length is greater than the maximum LBA, then the device:

a) shall return command aborted; and

b) may evict LBA ranges that are valid

Byte	7	6	5	4	3	2	1	0
0	(MSB)							
	LBA Range 0 (see [Editor's note Figure 250+3]Figure)							
7								(LSB)

8	(MSB)
	LBA Range 1
15	(LSB)
	. . .
504	(MSB)
	LBA Range 63
511	(LSB)

Figure 250+2 - Output data from the host for the HYBRID EVICT command

Byte	7	6	5	4	3	2	1	0
0	(MSB)							
	Starting LBA							
5								(LSB)
6	(MSB)			Range Length				
7								(LSB)

Figure 250+3 - LBA Range

13.6.14 <13.6.8>[Editor's note 13.6.9]First-party DMA HBA Support (Informative)

[editors note: make no changes]

13.7 [Editor's note 13.7]SATA Logs

[editors note: make the following changes to log addresses]

13.7.5 <13.7.1>[Editor's note 13.7.2]Log Address Overview

There are several log "files" available in a SATA device. They are all read-only. The READ LOG EXT command in the General Purpose Logging feature set (see ATA8-ACS) are used to read the SATA logs. In some cases the READ LOG DMA EXT command may also be used (see ACS-2).

Each log has an "address" by which it is referenced. Each log contains zero or ~~of~~ more "pages" of data. Each "page" contains 512 bytes of data.

13.7.1 <13.7.1>[Editor's note 13.7.2] Log Address Definitions

The log addresses assigned for Serial ATA are defined in Table 94.

Table 93 – Log Addresses for Serial ATA

Log Address	Description	Reference
00h~~ -~~ ..0Fh	As defined in the ATA8-ACS standard	
10h	NCQ Queued Error log	13.7.7
11h	Phy Event Counters log	13.7.8
12h	NCQ NON-DATA ~~Queue Management~~ log	13.7.9
13h	NCQ Send and Receive log	13.7.10
14h	Hybrid Information Log	13.7.12
15h~~14h -~~ ..17h	Reserved	
18h~~—~~ ..FFh	As defined in the ATA8-ACS standard	

13.7.6 <13.7.2>[Editor's note 13.7.3]General Purpose Log Directory (00h)

[editors note: make the following changes to the GPL log]

Devices supporting the Queued Error Log (see 13.7.7) reflect this support in the General Purpose Log Directory (see Table 95) log (00h) by having the value 1 at offset 20h and the value 0 at offset 21h of that log to indicate existence of a log at address 10h of 1 page in length.

Devices supporting the Phy Event Counters Log reflect this support in the General Purpose Log Directory (00h) by having the value 1 at offset 22h and the value 0 at offset 23h of that log to indicate existence of a log at address 11h of 1 page in length.

Devices supporting the NCQ NON-DATA ~~Queue Management~~ Log reflect this support in the General Purpose Log Directory (log 00h) by having the value 1 at offset 24h and the value 0 at offset 25h of that log to indicate existence of a log at address 12h of 1 page in length.

Devices supporting the NCQ Send and Receive Log reflect this support in the General Purpose Log Directory (00h) by having the value 1 at offset 26h and the value 0 at offset 27h of that log to indicate existence of a log at address 13h of 1 page in length.

Devices supporting the Hybrid Information Log reflect this support in the General Purpose Log Directory (00h) by having the value 1 at offset 28h and the value 0 at offset 29h of that log to indicate existence of a log at address 14h of 1 page in length.

Table 94 – General Purpose Log directory values for Serial ATA

Byte	Value
0~~–~~..1Fh	As defined in the ATA8-ACS standard
20h	1 if Native Command Queuing is supported, 0 if Native Command Queuing is not supported
21h	0
22h	1 if Phy Event Counters are supported 0 if Phy Event Counters are not supported
23h	0
24h	1 if NCQ NON-DATA log ~~Queue Management~~ is supported 0 if NCQ NON-DATA log ~~Queue Management~~ is not supported
25h	0
26h	1 if NCQ Send and Receive log is supported 0 if NCQ Send and Receive log is not supported
27h	0
28h	1 if Hybrid Information log is supported 0 if Hybrid Information log is not supported
29h	0
2Ah~~28h –~~..2Fh	Reserved
30h~~–~~..1FFh	As defined in the ATA8-ACS standard

13.7.7 ~~<13.7.3>~~[Editor's note 13.7.4]Queued Error Log (10h)

[editors note: make no changes]

13.7.8 ~~<13.7.4>~~[Editor's note 13.7.5]Phy Event Counters Log (11h)

[editors note: make no changes]

13.7.9 ~~<13.7.5>~~[Editor's note 13.7.6]NCQ NON-DATA ~~Queue Management~~ Log (12h)

[editors note: make the following changes to the NCQ queue management log]

13.7.9.4 [Editor's note 13.7.6.1] NCQ NON_DATA Log Overview

To determine the supported ~~NCQ QUEUE MANAGEMENT~~ NCQ NON-DATA subcommands and their respective features, host software ~~shall~~ may read log 12h.

This log shall be supported if the ~~NCQ QUEUE MANAGEMENT~~ NCQ NON-DATA command is supported (i.e., IDENTIFY DEVICE word 77 bit 5 is set to one.)

Dword	Bits	Description
0	Subcommand 0h	
	31~~-~~..5	Reserved ~~for NCQ Abort features~~
	4	Supports ~~4h Abort Type (~~Abort Selected TTag~~)~~
	3	Supports ~~3h Abort Type (~~Abort Non-Streaming~~)~~
	2	Supports~~ 2h Abort Type (~~Abort Streaming~~)~~
	1	Supports ~~1h Abort Type (~~Abort All~~)~~
	0	Supports Abort NCQ ~~(subcommand 0h)~~
1	Subcommand 1h	
	31~~-~~..3	Reserved~~ for NCQ Deadline features~~
	2	Supports Read Data Not Continue
	1	Supports Write Data Not Continue
	0	Supports ~~Deadline Handling~~ DEADLINE HANDLING ~~(subcommand 1h)~~
2	Subcommand 2h	
	31..1	Reserved
	0	Supports HYBRID DEMOTE BY SIZE
3	Subcommand 3h	
	31..1	Reserved
	0	Supports HYBRID CHANGE BY LBA RANGE
4	Subcommand 4h	
	31..1	Reserved
	0	Supports HYBRID CONTROL
…	…	…
15	Subcommand Fh	
	31~~-~~..1	Reserved ~~for subcommand Fh features~~
	0	Supports subcommand Fh

~~128-16~~ 16..127	31-..0	Reserved

Figure 253 – NCQ ~~Queue Management~~ NON-DATA Log (12h) data structure definition

13.7.9.5 [Editor's note 13.7.6.2]Supports Abort NCQ

If the Supports Abort NCQ bit is set to one, then the device supports the ABORT NCQ QUEUE command (see 13.6.11.5). If the Supports Abort NCQ bit is cleared to zero, then the device does not support the ABORT NCQ QUEUE command.

13.7.9.6 [Editor's note 13.7.6.3]Supports Abort All

If the Supports Abort All bit is set to one, then the device supports the ABORT NCQ QUEUE command (see 13.6.11.5) with the value of 0h in the Abort Type field. If the Supports Abort All bit is cleared to zero, then the device does not support the ABORT NCQ QUEUE command with the value of 0h in the Abort Type field.

13.7.9.7 [Editor's note 13.7.6.4]Supports Abort Streaming

If the Supports Abort Streaming bit is set to one, then the ABORT NCQ QUEUE command (see 13.6.11.5) supports the value of 1h in the Abort Type field. If the Supports Abort All bit is cleared to zero, then the ABORT NCQ QUEUE command does not support the value of 1h in the Abort Type field.

13.7.9.8 [Editor's note 13.7.6.5]Supports Abort Non-Streaming

If the Supports Abort Non-Streaming bit is set to one, then the device supports the ABORT NCQ QUEUE command (see 13.6.11.5) with the value of 2h in the Abort Type field. If the Supports Abort All bit is cleared to zero, then the device does not support the ABORT NCQ QUEUE command with the value of 2h in the Abort Type field.

13.7.9.9 [Editor's note 13.7.6.6]Supports Abort Selected TTAG

If the Supports Abort Selected TTAG bit is set to one, then the device supports the ABORT NCQ QUEUE command (see 13.6.11.5) with the value of 3h in the Abort Type field. If the Supports Abort All bit is cleared to zero, then the device does not support the ABORT NCQ QUEUE command with the value of 3h in the Abort Type field.

13.7.9.10 [Editor's note 13.7.6.7]Supports ~~Deadline Handling~~ DEADLINE HANDLING

If the Supports ~~Deadline Handling~~ DEADLINE HANDLING bit is set to one, then the device supports the DEADLINE HANDLING command (see 13.6.5.1). If the Supports ~~Deadline Handling~~ DEADLINE HANDLING bit is cleared to zero, then the device does not support the DEADLINE HANDLING command.

13.7.9.11 [Editor's note 13.7.6.8]Supports Write Data Not Continue

If the Supports Write Data Not Continue bit is set to one, then the device supports the WDNC field of the DEADLINE HANDLING command (see 13.6.5.1). If the Supports Write Data Not Continue bit is cleared to zero, then the device does not support the WDNC field of the DEADLINE HANDLING command (see 13.6.5.1).

13.7.9.12 [Editor's note 13.7.6.9]Supports Read Data Not Continue

If the Supports Read Data Not Continue bit is set to one, then the device supports the RDNC field of the DEADLINE HANDLING command (see 13.6.5.1). If the Supports Read Data Not Continue bit is cleared to zero, then the device does not support the RDNC field of the DEADLINE HANDLING command (see 13.6.5.1xxx).

13.7.9.13 [Editor's note 13.7.6.10]Supports HYBRID DEMOTE BY SIZE

If the Supports HYBRID DEMOTE BY SIZE bit is set to one, then the device supports the HYBRID DEMOTE BY SIZE command (see 13.6.5.4). If the Supports HYBRID DEMOTE BY SIZE bit is cleared to zero, then the device does not support the HYBRID DEMOTE BY SIZE command.

13.7.9.14 [Editor's note 13.7.6.11]Supports HYBRID CHANGE BY LBA RANGE

If the Supports HYBRID CHANGE BY LBA RANGE is set to one, then the device supports the HYBRID CHANGE BY LBA RANGE command (see13.6.5.5). If the Supports HYBRID CHANGE BY LBA RANGE bit is cleared to zero, then the device does not support the HYBRID CHANGE BY LBA RANGE command.

13.7.9.15 [Editor's note 13.7.6.12]Supports HYBRID CONTROL

If the Supports HYBRID CONTROL bit is set to one, then the device supports the HYBRID CONTROL command (see 13.6.5.6). If the Supports HYBRID CONTROL bit is cleared to zero, then the device does not support the HYBRID CONTROL command.

13.7.10 ~~<13.7.6>~~[Editor's note 13.7.7]NCQ Send and Receive Log (13h)

13.7.10.4 ~~<13.7.6>~~[Editor's note 13.7.7.1]NCQ Send and Receive Log overview

[editors note: make the following changes to the NCQ Send and Receive log]

To determine the supported SEND FPDMA QUEUED and RECEIVE FPDMA QUEUED subcommands and their respective features, host software ~~shall~~ may read log 13h.

This ~~page~~ log shall be supported if the SEND FPDMA QUEUED and RECEIVE FPDMA QUEUED command is supported (IDENTIFY DEVICE word 77 bit 6 is set to one.)

Dword	Bits	Description
0	Subcommands Supported	
~~0~~	31..~~1~~2	Reserved
	1	Supports HYBRID EVICT (see 13.7.10.5)
	0	~~1 =~ Supports ~~the SEND FPDMA QUEUED~~ Data Set Management ~~subcommand~~ (see 13.7.10.6) ~~0 = Does not support the SEND FPDMA QUEUED Data Set Management subcommand~~
~~1~~	Data Set Management	
1	31~~-~~..1	Reserved~~ for future SEND FPDMA QUEUED Data Set Management subcommand supported attributes~~
	0	~~1 = The SEND FPDMA QUEUED Data Set Management subcommand supports the Trim Attribute~~Supports Trim (see 13.7.10.7) ~~0 = The SEND FPDMA QUEUED Data Set Management subcommand does not support the Trim Attribute~~
~~128-232~~..127	~~31-0~~	Reserved

Figure 254 – Log ~~Page~~ 13h – NCQ SEND and RECEIVE Log

13.7.10.5 [Editor's note 13.7.7.2]Supports HYBRID EVICT

If the Supports HYBRID EVICT bit is set to one, then the device supports the HYBRID EVICT subcommand (see 13.6.13.10) of the SEND FPDMA QUEUED command. If the Supports HYBRID EVICT bit is cleared to zero, then the device does not support the HYBRID EVICT subcommand of the SEND FPDMA QUEUED command.

13.7.10.6 [Editor's note 13.7.7.3]Supports Data Set Management

If the Supports Data Set Management bit is set to one, then the device supports the Data Set Management subcommand (see 13.6.13.9) of the SEND FPDMA QUEUED command. If the Supports Data Set Management bit is cleared to zero, then the device does not support the Data Set Management subcommand of the SEND FPDMA QUEUED command.

13.7.10.7 [Editor's note 13.7.7.4]Supports Trim

If the Supports Trim bit is set to one, then the device supports the Trim attribute of the Data Set Management subcommand (see 13.6.13.9) of the SEND FPDMA QUEUED command. If the Supports Data Set Management bit is cleared to zero, then the device does not support the Trim attribute of the Data Set Management subcommand of the SEND FPDMA QUEUED command.

13.7.11 ~~<13.7.7>~~[Editor's note 13.7.8]Identify Device Data Log (30h)

[editors note: make no changes to the Identify Device Data log]

[editors note: add this new section 13.7.12. The editor can rearrange later if needed.]

13.7.12 [Editor's note 13.7.9]Hybrid Information Log (14h)

13.7.12.4 [Editor's note 13.7.9.1]Overview

If the Hybrid Information feature is supported, then the Hybrid Information log shall be supported. The Hybrid Information log consists of one page (see [Editor's note Figure 254+1]). The log is read-only. Reading the log shall not cause the device to change power management state.

Byte	Description	Reference
0	Hybrid Information Header	13.7.12.5
64	Hybrid Information Descriptor for Hybrid Priority 0	13.7.12.6
...	...	
	Hybrid Information Descriptor for Maximum Priority (N)	13.7.12.6
64 + (16 × N) to 511	Padding	

Figure 254+1 - Hybrid Information log data

Padding
Data transfer lengths shall be non-zero multiples of 512 bytes. Pad bytes shall be appended as needed to meet this requirement. Pad bytes shall have a value of 00h.

N is the number of Hybrid Information Descriptors reported in the log.

13.7.12.5 [Editor's note 13.7.9.2]Hybrid Information Header

[editors note: the subclause numbers are messed up]

13.7.5.4.1 [Editor's note 13.7.9.2.1]Overview

[Editor's note Figure 254+2] describes the Hybrid Information Header which contains summary information for the hybrid device.

Byte	Type	Description	Reference
0..1	Word	Bits Description 15:4 Reserved	13.7.5.4.2

Byte	Type	Description	Reference
		3:0 Number of Hybrid Information Descriptors	
2	Byte	Enabled	13.7.5.4.3
3	Byte	Hybrid Health	13.7.5.4.4
4	Byte	Dirty Low Threshold	13.7.5.4.4
5	Byte	Dirty High Threshold	13.7.5.4.6
6	Byte	Optimal Write Granularity	13.7.5.4.7
7	Byte	Bits Description 7:4 Reserved 3:0 Maximum Hybrid Priority Level	13.7.5.4.8
8	Byte	Power Condition	13.7.5.4.9
9	Byte	Caching Medium Enabled	13.7.5.4.10
10	Byte	Supported Options	13.7.5.4.11
11		Reserved	
12..15	DWord	Time Since Enabled	13.7.5.4.12
16..23	QWord	NVM Size	[editors note: 13.7.5.4.13]
24..31	QWord	Enable Count	13.7.5.4.14
32..33	Word	Bits Description 7:5 Reserved 4:0 Maximum Eviction Commands	13.7.5.4.15
34..35	Word	Maximum Eviction Data Blocks	13.7.5.4.16
36..63		Reserved	

Figure 254+2 - Hybrid Information Header

13.7.5.4.2 [Editor's note 13.7.9.2.2]Number of Hybrid Information Descriptors

The Number of Hybrid Information Descriptors field indicates the the number of Hybrid Information descriptors that follow the header.

13.7.5.4.3 [Editor's note 13.7.9.2.3]Enabled

Table 96 indicates the value that the device shall indicate in IDENTIFY DEVICE in each case.

Table 95 - Hybrid Information Enabled Field

Value	Description	IDENTIFY DEVICE data word 79 bit 9
00h	Hybrid Information Disabled	0
80h	Hybrid Information Disable In Process	0
FFh	Hybrid Information Enabled	1
All other values	Reserved	

13.7.5.4.4 [Editor's note 13.7.9.2.4]Hybrid Health

13.7.5.5.4.2 [Editor's note 13.7.9.2.4.1]Overview

Table 95+2] describes the values of the Hybrid Health field, which contains several indicators of the health of the non-volatile caching medium.

NOTE: If the non-volatile caching medium is healthy, the Hybrid Health field is cleared to zero.

Table 95+2 - Hybrid Health

Bit	Description	Reference
7:4	Reserved	
3	Data Loss	[Editor's note 13.7.9.2.4.2]
2	Read Only	[Editor's note 13.7.9.2.4.3]
1	NVM Size Changed	[Editor's note 13.7.9.2.4.4]
0	Unuseable	[Editor's note 13.7.9.2.4.5]

13.7.5.5.4.2 [Editor's note 13.7.9.2.4.2]Data Loss

If the Data Loss bit is set to one, then some of the data in the non-volatile caching medium has become inaccessible since the Hybrid Information log was last read (see 13.20.9.6).

If the Data Loss bit is cleared to zero, then no data loss has been detected in the non-volatile caching medium since the Hybrid Information log was last read.

13.7.5.5.4.2 [Editor's note 13.7.9.2.4.3]Read Only

If the Read Only bit is set to one, then the non-volatile caching medium is read only (see 13.20.9.5).

If the Read Only bit is cleared to zero, then the non-volatile caching medium may be read or written.

13.7.5.5.4.2 [Editor's note 13.7.9.2.4.4]NVM Size Changed

If the NVM Size Changed bit is set to one, then the device has changed the NVM Size of the non-volatile caching medium since the Hybrid Information log was last read (see 13.20.9.4).

If the NVM Size Changed bit is cleared to zero, then the device has not changed the NVM Size of the non-volatile caching medium since the Hybrid Information log was last read.

13.7.5.5.4.2 [Editor's note 13.7.9.2.4.5]Unuseable

If the Unuseable bit is set to one, then the non-volatile caching medium is no longer useable (see 13.20.9.7).

If the Unuseable bit is cleared to zero, then the non-volatile caching medium is useable.

13.7.5.4.5 [Editor's note 13.7.9.2.5]Dirty Low Threshold

The Dirty Low Threshold field indicates the threshold for the amount of dirty user logical sectors in the non-volatile caching medium at which point syncing operations should stop. For additional information, see the HYBRID CONTROL command (see 13.6.5.6).

13.7.5.4.6 [Editor's note 13.7.9.2.6]Dirty High Threshold

The Dirty High Threshold field indicates the threshold for the amount of dirty user logical sectors in the non-volatile caching medium at which point syncing operations should begin. For additional information, see the HYBRID CONTROL command (see 13.6.5.6).

13.7.5.4.7 [Editor's note 13.7.9.2.7]Optimal Write Granularity

The Optimal Write Granularity field indicates the optimal number of logical sectors for the host to write to the non-volatile caching medium, expressed as a power of 2. If the field contains FFh, then the optimal write granularity is not indicated.

> For example: 0 indicates 2^0 = 1 logical sector, 1 indicates 2^1 = 2 logical sectors, 2 indicates 2^2 = 4 logical sectors.

13.7.5.4.8 [Editor's note 13.7.9.2.8]Maximum Hybrid Priority Level

The Maximum Hybrid Priority Level field indicates the maximum supported value of the Hybrid Priority field (see 13.20.6.5). The Maximum Hybrid Priority Level shall be non-zero.

13.7.5.4.9 [Editor's note 13.7.9.2.9]Power Condition

The Power Condition field indicates the current power condition as the CHECK POWER MODE command would report in normal outputs (see ACS-2).

13.7.5.4.10 [Editor's note 13.7.9.2.10]Caching Medium Enabled

The Caching Medium Enabled field indicates whether or not the non-volatile caching medium is usable by the host or the device (see 13.6.5.6).

If the Caching Medium Enabled field is set to FFh, then the non-volatile caching medium is enabled for use by the host and the device.

If the Caching Medium Enabled field is cleared to 00h, then the non-volatile caching medium shall not be used by either the host or the device .

13.7.5.4.11 [Editor's note 13.7.9.2.11]Supported Options

The Supported Options field (see [Editor's note Table 95+3]) indicates which optional behavior is supported.

Table 95+3 - Supported Options

Bit	Description
7:1	Reserved
0	Max Priority Behavior

If the Max Priority Behavior bit is set to one (see 13.20.5.5.2.2), then the device shall:

a) insert logical sectors into the non-volatile caching medium if a command specifies the maximum hybrid priority level;

b) abort any command that specifies the maximum hybrid priority level if there are not enough available logical sectors in the non-volatile caching medium;

c) support the HYBRID EVICT command (see 13.6.13.10); and

d) support the HYBRID CHANGE BY LBA RANGE command (see 13.6.5.5).

If the Max Priority Behavior bit is cleared to zero (see 13.20.5.5.2.2), then the device should insert logical sectors into the non-volatile caching medium when a command specifies the maximum hybrid priority level.

13.7.5.4.12 [Editor's note 13.7.9.2.12]Time Since Enabled

The Time Since Enabled field indicates the number of power-on hours since the Hybrid Information feature was enabled. This is an unsigned integer.

13.7.5.4.13 [Editor's note 13.7.9.2.13]NVM Size

The NVM Size field indicates the number of logical sectors that comprise the non-volatile caching medium.

Note: The value of the NVM Size field may vary over time because of vendor specific factors.

13.7.5.4.14 [Editor's note 13.7.9.2.14]Enable Count

The Enable Count field contains an unsigned integer incremented by one by the device as a result of succesfully enabling the Hybrid Information feature (see 13.3.15).

13.7.5.4.15 [Editor's note 13.7.9.2.15]Maximum Eviction Commands

The Maximum Eviction Commands field indicates the maximum number of HYBRID EVICT commands (see 13.6.13.10) that the device supports in the command queue at the same time. A value of zero indicates that the device does not limit the number of HYBRID EVICT commands in the queue.

13.7.5.4.16 [Editor's note 13.7.9.2.16]Maximum Eviction Data Blocks

The Maximum Eviction Data Blocks field limits the maximum number of data blocks that may be specified in the HYBRID EVICT command (see 13.6.13.10).

13.7.12.6 [Editor's note 13.7.9.3]Hybrid Information Descriptor

13.7.5.4.1 [Editor's note 13.7.9.3.1]Overview

Figure [Editor's note Figure 254+3] describes the Hybrid Information Descriptor. There shall be one Hybrid Information Descriptor returned for each supported Hybrid Priority value, in order of increasing Hybrid Priority Level.

Byte	Type	Description	Reference
0	Byte	Hybrid Priority	13.7.5.4.2
1	Byte	Consumed NVM Size Fraction	13.7.5.4.3
2	Byte	Consumed Mapping Resources Fraction	13.7.5.4.4
3	Byte	Consumed NVM Size For Dirty Data Fraction	13.7.5.4.5
4	Byte	Consumed Mapping Resources For Dirty Data Fraction	13.7.5.4.6
5..15		Reserved	

Figure 254+3 - Hybrid Information Descriptor

13.7.5.4.2 [Editor's note 13.7.9.3.2]Hybrid Priority

The Hybrid Priority field indicates the Hybrid Priority number that this descriptor represents..

13.7.5.4.3 [Editor's note 13.7.9.3.3]Consumed NVM Size Fraction

The value of the Consumed NVM Size Fraction field, when divided by 255, indicates the fraction (see [Editor's note Figure 254+4]) of the NVM Size for this Hybrid Priority's logical sectors that is currently consumed (i.e., used). The value is an unsigned integer from 00h to FFh. The value 00h indicates that no NVM Size is currently consumed. The value FFh indicates that all of the NVM Size is currently consumed.

$$\text{Consumed Capacity Fraction} = \frac{A \times 255}{B}$$

where:

A = current number of logical sectors associated with this hybrid priority

B = NVM Size (see 13.7.5.4.13)

Figure 254+4 - Consumed NVM Size

13.7.5.4.4 [Editor's note 13.7.9.3.4]Consumed Mapping Resources Fraction

The value of the Consumed Mapping Resources Fraction field divided by 255 indicates the fraction of the mapping resources for this Hybrid Priority's logical sectors in the non-volatile caching medium that are currently consumed. The value is an unsigned integer from 00h to FFh. The value 00h indicates that no mapping resources are currently consumed. The value FFh indicates that all of the mapping resources are currently consumed.

13.7.5.4.5 [Editor's note 13.7.9.3.5]Consumed NVM Size For Dirty Data Fraction

The value of the Consumed NVM Size For Dirty Data Fraction field divided by 255 indicates the fraction (see [Editor's note Figure 254+5]) of the maximum NVM Size for this Hybrid Priority's data that is currently marked as dirty data. The value is an unsigned integer from 00h to FFh. The value 00h indicates that no NVM Size is currently consumed. The value FFh indicates that all of the NVM Size is currently consumed.

$$\text{Consumed Capacity for Dirty Data Fraction} = \frac{A \times 255}{B}$$

where:

A = current NVM Size consumed by dirty data associated with this Hybrid Priority level

B = NVM Size (see 13.7.5.4.13)

Figure 254+5 - Consumed NVM Size For Dirty Data

13.7.5.4.6 [Editor's note 13.7.9.3.6]Consumed Mapping Resources for Dirty Data Fraction

The value of the Consumed Mapping Resources For Dirty Data Fraction field divided by 255 indicates the fraction of the mapping resources for this Hybrid Priority's data in the non-volatile caching medium that are currently consumed for mapping dirty data. The value is an unsigned integer from 00h to FFh. The value 00h indicates that no mapping resources are currently consumed that relate to dirty data. The value FFh indicates that all of the mapping resources are currently consumed that relate to dirty data.

13.8 [Editor's note 13.8]Asynchronous Notification (Optional)

[editors note: make no changes]

13.9 [Editor's note 13.9]Phy Event Counters (Optional)

[editors note: make no changes]

13.10 [Editor's note 13.10]Hardware Feature Control (optional)

[editors note: make no changes]

13.11 [Editor's note 13.11]Staggered Spin-up (Optional)

[editors note: make no changes]

13.12 [Editor's note 13.12]Non-512 Byte Sector Size (Informative)

[editors note: make no changes]

13.13 [Editor's note 13.13]Defect Management (Informative)

[editors note: make no changes]

13.14 [Editor's note 13.14]Enclosure Services/Management (Optional)

[editors note: make no changes]

13.15 [Editor's note 13.15]HDD Activity Intication (Optional)

[editors note: make no changes]

13.16 [Editor's note 13.16]Port Multiplier Discovery and Enumeration

[editors note: make no changes]

13.17 [Editor's note 13.17]Automatic Partial to Slumber Transitions

[editors note: make no changes]

13.18 [Editor's note 13.18]Serial ATA Link Power Management Support

[editors note: make no changes]

13.19 [Editor's note 13.19]DHU Specific Operation (optional)

[editors note: make no changes]

[editors note: add new section 'Hybrid Information (Optional)' to clause 13]

13.20 [Editor's note 13.20] Hybrid Information Feature (Optional)

13.20.5 [Editor's note 13.20.1]Overview

A Solid State Hybrid Device (SSHD) contains both primary medium (e.g., rotating magnetic) and a non-volatile caching medium (e.g., flash memory).

The Hybrid Information feature allows the host to provide information (e.g., hints) to the device, which the device uses for various purposes (e.g., to decide which medium to save the user data to for optimal performance in retrieval).

SSHDs may determine data to cache based on observed accesses of LBAs and length of incoming requests. However, SSHDs do not possess host information to make the most optimal caching decisions (e.g., file type associated with the LBAs in the request). The Hybrid Information feature allows the host to be involved in making caching decisions to aid in optimizing cache utilization.

This feature provides a method for the host to indicate the caching priority of incoming requests to the device, and feedback to the host on how much non-volatile caching medium has been consumed at the various caching priority levels.

The use of the maximum hybrid priority level is determined by the MaxPriority Behavior bit (see 13.7.5.4.11) in the Hybrid Information log.

If the MaxPriority Behavior bit is set to one, then the highest caching priority passed to the device instructs the SSHD that this data shall be placed and remain in the non-volatile caching medium until explicitly evicted by the host.

If the MaxPriority Behavior bit is cleared to zero, then the highest caching priority passed to the device instructs the SSHD that this data should:

a) be placed in non-volatile caching medium; and

b) remain in the non-volatile caching medium.

Intermediate caching priority levels inform the device of the importance of the data being placed in the non-volatile caching medium, but makes no requirement for the device to place the logical sectors in the non-volatile caching medium (e.g., the host only imparts a relative caching level in relation to other requests). The device should make the best decision possible based on the caching priority provided by the host and other device knowledge (e.g. rotational position optimizations).

Having multiple caching priorities allows for the host to group data by importance. Data that is required to be in the non-volatile caching medium is passed with the highest priority level . The intermediate priorities may be used for data that would provide value to the user if present in the non-volatile caching medium , but is not required to meet power budgets or responsiveness criterion (e.g. medium files, application data, etc).

In order to maintain a given responsiveness for insertions into the non-volatile caching medium, the host may specify high and low dirty thresholds to ensure that the device does not consume too much internal bandwidth syncing logical sectors between the non-volatile caching medium and the primary medium, while still leaving room to absorb new writes.

The hints are attached to some commands sent by the host, included in the Register Host to Device FIS (see [editors note: 10.3.4] and 13.20.6).

The host may evict logical sectors from the non-volatile caching medium using the HYBRID EVICT command (see 13.6.13.10).

If the Hybrid Information feature is not supported (i.e., IDENTIFY DEVICE word 78 bit 9 is cleared to zero), then:

a) the device shall ignore hints as described in this subclause for any command; and

b) the device shall not support:
   A) the SET FEATURES Enable/Disable Hybrid Information subcommand (see 13.3.15);
   B) the Hybrid Information log (see 13.7.12);
   C) the HYBRID EVICT command (see 13.6.13.10);
   D) the HYBRID DEMOTE BY SIZE command (see13.6.5.4);
   E) the HYBRID CHANGE BY LBA RANGE command (see 13.6.5.5); and
   F) the HYBRID CONTROL command (see 13.6.5.6).

If the Hybrid Information feature is supported (i.e., IDENTIFY DEVICE word 78 bit 9 is set to one), then:

a) the device shall support:
   A) the following feature sets:
      a)NCQ Autosense (i.e., IDENTIFY DEVICE data word 78 bit 7 shall be set to one); and
      b)POWER UP IN STANDBY feature set (see ACS-2);
   B) the following SET FEATURES subcommands:
      a)Enable/disable the POWER UP IN STANDBY feature set (see ACS-2); and
      b)Enable/Disable Hybrid Information (see 13.3.15)
   C) the HYBRID DEMOTE BY SIZE command (see13.6.5.4); and
   D) the HYBRID CONTROL command (see 13.6.5.6); and
   E) the following logs:
      a)the NCQ NON-DATA log (see 13.7.9);
      b)the NCQ Send and Receive log (see 13.7.10);
      c)the Identify Device Data log (see 13.7.11); and
      d)the Hybrid Information log (see 13.7.12);

b) the device should support the Device Sleep feature (i.e., IDENTIFY DEVICE data word 78 bit 8 should be set to one);

c) the device shall not support the SET FEATURES POWER UP IN STANDBY feature set device spin-up subcommand (see ACS-2);

d) the following commands shall be able to return command completion without error while in the PM2:Standby state (see ACS-3):

   A) all read commands, if the requested logical sectors are in the non-volatile caching medium;

   B) all write commands, if the device stores all of the data for the command in the non-volatile caching medium;

   C) commands to read the following logs:

      a) Identify Device Data log;

      b) Hybrid Information log;

      c) Power Conditions log;

      d) Queued Error log;

      e) NCQ NON-DATA log; and

      f) NCQ Send and Receive log;

   D) IDENTIFY DEVICE command;

   E) CHECK POWER MODE command;

   F) SMART RETURN STATUS command; and

   G) SECURITY UNLOCK command, if supported;

e) if the Hybrid Information feature is enabled (i.e., IDENTIFY DEVICE word 79 bit 9 is set to one), then:

   A) the device shall process the Hybrid Information field as described in this subclause for:

      a) the READ FPDMA QUEUED command (see 13.6.8);
      b) the WRITE FPDMA QUEUED command (see 13.6.9);
      c)the HYBRID DEMOTE BY SIZE command (see13.6.5.4);
      d)the HYBRID CHANGE BY LBA RANGE command (see 13.6.5.5);
      e)the HYBRID CONTROL command (see 13.6.5.6); and
      f) these non-NCQ commands:
         A) the READ DMA EXT command;
         B) the WRITE DMA EXT command; and
         C) the WRITE DMA FUA EXT command;

   and

   B) for all other commands, the device shall ignore the Hybrid Information field as described in this subclause;

and

f) if the Hybrid Information feature is disabled (i.e., IDENTIFY DEVICE word 79 bit 9 is cleared to zero), then the device shall ignore the Hybrid Information field as described in this subclause.

13.20.6 [Editor's note 13.20.2]Hybrid Information Field Bits

13.20.6.4 [Editor's note 13.20.2.1]Hybrid Information Field Bits overview

[Editor's note Figure 279+1] describes the mapping of the Hybrid Information field, which is transported in the Auxiliary(23:16) field of the Register Host to Device FIS (see 10.3.9)

If the Hybrid Information feature is not supported, then the device shall ignore the Hybrid Priority field.

If:

a) the Hybrid Information feature is supported;
b) the Hybrid Information feature is enabled; and
c) the Hybrid Information Is Valid bit is set to one,

then the Hybrid Priority field is valid.

If:

a) the Hybrid Information feature is supported; and
b) the Hybrid Information feature is not enabled or the Hybrid Information Is Valid field is cleared to zero,

then the device shall ignore the Hybrid Priority field.

Auxiliary Field Bit	Description	Reference
16-19	Hybrid Priority	13.20.6.5
20	Reserved	
21	Hybrid Information Is Valid	
22-23	Reserved	

Figure 279+1 - Hybrid Information Field

13.20.6.5 [Editor's note 13.20.2.2]Hybrid Priority

13.20.5.4.1 [Editor's note 13.20.2.2.1]Hybrid Priority overview

The Hybrid Priority field indicates the hybrid priority to be associated with the logical sectors when processing a command for which the Hybrid Information fields are valid (see 13.20.6).

If a command specifies a Hybrid Priority value which is greater than the value of the Maximum Hybrid Priority Level field (see 13.7.5.4.8), then the device shall return command aborted.

If the device processes a read command, then for each logical sector the device should read from the medium with the shortest latency that contains the most recent copy of the requested logical sectors.

13.20.5.4.2 [Editor's note 13.20.2.2.2]Hybrid Priority field equal to Maximum Hybrid Priority Level

13.20.5.5.2.2 [Editor's note 13.20.2.2.2.1]Hybrid Priority field equal to Maximum Hybrid Priority Level overview

The behavior of the highest permitted Hybrid Priority value (i.e., Maximum Hybrid Priority Level) is specified by the value of the Max Priority Behavior bit of the Supported Options field (see 13.7.5.4.11) in the Hybrid Information log.

13.20.5.5.2.2 [EDITOR'S NOTE 13.20.2.2.2.2]Max Priority Behavior set to one

If:

a) the device processes a read command or a write command, with the Hybrid Information Is Valid field set to one and the Hybrid Priority value set equal to the Maximum Hybrid Priority Level; and

b) the Max Priority Behavior bit is set to one,

then:

a) the device shall not evict logical sectors associated with the Maximum Hybrid Priority Level from the non-volatile caching medium in order to insert the logical sectors specified by the command;

b) the device may evict logical sectors associated with a lower hybrid priority from the non-volatile caching medium in order to insert the logical sectors specified by the command;

c) if:

   A. the logical sectors specified by the command are not currently in the non-volatile caching medium; and

   B. the device does not have capacity available for all of the requested logical sectors in the non-volatile caching medium,

   then the device shall:

   A. return command aborted; and

   B. in the Queued Error log:

      a. the Sense Key field shall be set to COMMAND ABORTED; and

b. the additional sense code (i.e., ASC field and ASCQ field) set to INSUFFICIENT RESOURCES;

d) copy the requested logical sectors to the non-volatile caching medium if the logical sectors specified by the command are not currently in the non-volatile caching medium and the non-volatile caching medium has capacity available for all of the requested logical sectors; and

e) associate the logical sectors specified by the command with the Hybrid Priority value specified by the command and return command complete with no error.

13.20.5.5.2.2 [EDITOR'S NOTE 13.20.2.2.2.3]Max Priority Behavior is zero

If:

a) the device processes a read command or a write command, with the Hybrid Information Is Valid set to one and the Hybrid Priority value set to the Maximum Hybrid Priority Level; and

b) the Max Priority Behavior bit is cleared to zero,

then:

a) the device may evict any logical sectors from the non-volatile caching medium in order to insert the logical sectors specified by the command;

b) the device should put the requested logical sectors in the non-volatile caching medium;

c) associate the logical sectors that were put in the non-volatile caching medium with the Hybrid Priority value specified by the command and return command complete with no error.

13.20.5.4.3 [Editor's note 13.20.2.2.3]Hybrid Priority field is less than Maximum Hybrid Priority Level and is greater than zero

If the device processes a read command or a write command, with the non-zero Hybrid Priority value set less than the Maximum Hybrid Priority Level, then:

a) the device may evict logical sectors from the non-volatile caching medium in order to insert the logical sectors specified by the command with Hybrid Priority value if the logical sectors in the non-volatile caching medium:

A. have a lower Hybrid Priority; or

B. have the same Hybrid Priority but are older;

b) if the logical sectors specified by the command are not currently in the non-volatile caching medium, then the device:

A. should put the requested logical sectors to the non-volatile caching medium; and

B. may put the requested logical sectors to the non-volatile caching medium after returning command completion;

and

c) associate the logical sectors specified by the command with the Hybrid Priority value specified by the command and return command completewith no error.

13.20.5.4.4 [Editor's note 13.20.2.2.3]Hybrid priority is zero

If the device processes a read command or a write command, with the Hybrid Priority value cleared to zero, then:

a) no Hybrid Priority preference is specified;

b) the device should write to primary medium; and

c) the device should not put the selected logical sectors in the non-volatile caching medium.

13.20.7 [Editor's note 13.20.3]Syncing

The Hybrid Information log indicates several parameters related to syncing:

a) Dirty Low Threshold (see 13.7.5.4.5); and

b) Dirty High Threshold (see 13.7.5.4.6).

The device should prioritize the logical sectors contained in the non-volatile caching medium to be synched as follows:

1. starting from the Hybrid Priority zero, up to the Maximum Hybrid Priority Level;
2. within each Hybrid Priority, sync the least recently used logical sectors first and then the most recently used logical sectors.

13.20.8 [Editor's note 13.20.4]Interactions with ATA Power Management

If :

a) The Hybrid Information feature is enabled;

b) the device processes a read command or a write command; and

c) the requested logical sectors are not accessible in the current power condition,

then:

a) the device shall return command aborted; and

b) if the failing command is READ FPDMA QUEUED or WRITE FPDMA QUEUED, then the device shall set the following in the Queued Error log:

   a. the Sense Key field shall be set to NOT READY; and

   b. the additional sense code (i.e., ASC field and ASCQ field) shall be set to LOGICAL UNIT NOT READY – INITIALIZING COMMAND REQUIRED.

If the device indicates a sense key of NOT READY and an additional sense code of LOGICAL UNIT NOT READY – INITIALIZING COMMAND REQUIRED, then the device should be explicitly spun up before the host reissues the command.

If the power condition was entered as a result of processing the EPC Go To Power Condition command (see ACS-3) with the Hold Power Condition bit cleared to zero, then see [Editor's note Table 96+1] for interactions with the current power source reported by the device.

Table 96+1 - Current Power Source Interactions

CURRENT POWER SOURCE[a]	**Description**
0h	vendor specific
1h	The device should not go to a power condition that consumes more power to sync logical sectors
2h	Syncing operations should not adversely affect performance. The device should be more aggressive about syncing than when on battery. If the device changed power condition to process syncing, then the device should return to the previous power condition on completion of the syncing operation.
[a] The CURRENT POWER SOURCE field is indicated in the IDENTIFY DEVICE data log, page 04h	

13.20.9 [Editor's note 13.20.5]Other Hybrid Conditions

13.20.9.4 [Editor's note 13.20.5.1]NVM Size changed

The device may reduce the NVM Size of the non-volatile caching medium. If the device reduces the NVM Size of the non-volatile caching medium, then

a) the device shall set the NVM Size Changed bit to one in the Hybrid Health field (see 13.7.5.4.4) of the Hybrid Information log; and

b) if the Hybrid Information feature is enabled and the device processes a SMART RETURN STATUS command, then the device shall:

   A. set the value of LBA(23:8) to 2CF4h (i.e., the device has detected a threshold exceeded condition); and

   B. return command complete with no error.

If the host reads the Hybrid Information log, then the device shall clear the NVM Size Changed bit after returning the log data to the host.

13.20.9.5 [Editor's note 13.20.5.2]Read Only

The device may change the non-volatile caching medium to read-only access. If the device changes the non-volatile caching medium to read-only access, then:

a) the device shall set the Read Only bit to one in the Hybrid Health field (see 13.7.5.4.4) of the Hybrid Information log; and

The device may clear the Read Only bit under vendor specific conditions.

13.20.9.6 [Editor's note 13.20.5.3]Data Loss

If the device encounters conditions such that some logical sectors in the non-volatile caching medium are no longer accessible, then:

a) the device shall set the Data Loss bit to one in the Hybrid Health field (see 13.7.5.4.4) of the Hybrid Information log; and

b) if the Hybrid Information feature is enabled and the device processes a SMART RETURN STATUS command, then the device shall:

   A. set the value of LBA(23:8) to 2CF4h (i.e., the device has detected a threshold exceeded condition); and

   B. return command complete with no error.

If the host reads the Hybrid Information log, then the device shall clear the Data Loss bit after returning the log data to the host.

13.20.9.7 [Editor's note 13.20.5.4]Unuseable

If the device encounters conditions such that the non-volatile caching medium has become unuseable, then:

a) the device shall set the Unuseable bit to one in the Hybrid Health field (see 13.7.5.4.4) of the Hybrid Information log;

b) the device shall disable the Hybrid Information feature (i.e., IDENTIFY DEVICE word 79 bit 9 cleared to zero);

c) the device shall remove indication of support for the Hybrid Information feature (i.e., IDENTIFY DEVICE data word 79 bit 9 cleared to zero); and

d) if the device processes a SMART RETURN STATUS command, then the device shall:

   A. set the value of LBA(23:8) to 2CF4h (i.e., the device has detected a threshold exceeded condition); and

   B. return command complete with no error.

The device may clear the Unuseable bit under vendor specific conditions.

13.20.10 [Editor's note 13.20.6]Automatic Disable

If:

a) the Hybrid Information feature is currently enabled; and

b) the device has not processed any command to read the Hybrid Information log for 25 consecutive power cycles,

then the device shall:

1. change the Hybrid Priority for all logical sectors in the non-volatile caching medium to zero;
2. clear the Enabled field (see 13.6.5.3.3) in the Hybrid Information log to zero; and
3. disable the Hybrid Information feature (i.e., clear IDENTIFY DEVICE data word 79 bit 9 to zero).

14 [Editor’s note 14]Host adapter register interface

[editors note: make no changes]

15 [Editor’s note 15]Error handling

[editors note: make no changes]

16 [Editor’s note 16]Port Multiplier

[editors note: make no changes]

17 [Editor’s note 17]Port Selector

[editors note: make no changes]

