

**Proposed  
Draft**

**Serial ATA  
International Organization**

**Version 06  
April 29, 2013**

---

## **SATA31\_Hybrid Information Update**

**Title: Hybrid Information Update**

**Proposed change, new functionality, or behavior to Serial ATA Revision 3.1 and TPR042 – Hybrid Information Feature**

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

Permission is granted to the Promoters, Contributors and Adopters of the Serial ATA International Organization to reproduce this document for the purposes of evolving the technical content for internal use only without further permission provided this notice is included. All other rights are reserved and may be covered by one or more Non-Disclosure Agreements including the Serial ATA International Organization participant agreements. Any commercial or for-profit replication or republication is prohibited. Copyright © 2005 to 2013 Serial ATA International Organization. All rights reserved.

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

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

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

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

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

Copyright © 2005 to 2013 Serial ATA International Organization. All rights reserved.

## Author Information

Author Name	Company	e-mail address
Lee Prewitt	Microsoft	leprewit@microsoft.com

## Workgroup Chair Information

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

## Document History

Version	Date	Comments
00	01/14/2013	Initial draft
01	02/03/2013	Changed document number to D162. Added per command behavior indication.
02	02/10/2013	Clarified to not spin-up the drive. Updated CB definition.
03	03/18/2013	Added the full set of possible bit combinations to Table 96+3.
04	03/24/2013	Added cross references for the Cache Behavior Descriptions in Table 96+3.
05	03/25/2013	Member review
06	04/29/2013	Ratified

## Introduction

Systems need to determine the best policies to minimize power consumption. In order to accomplish this they instruct devices to enter 'almost off' states or completely remove power from the device itself. Because of this, the host needs to have good control over the spin state of the drive.

## Summary of the problem

Currently, if the hybrid device is spun down and the host issues a Change Priority by LBA command, depending on the state of the Max Priority Behavior bit, the new priority requested and the state of the cache, the drive will either be required to spin up and move the data into the cache or should do so. This is a problem if the platform is in a low power mode or if the platform is being transported by the user.

## Proposed changes

Add a new bit to the Supported Options table to indicate that the device supports per command tagging of the HYBRID CHANGE BY LBA RANGE command to indicate whether or not data should be pulled in to the cache if it is not already there.

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

Existing text includes SATA Rev 3.1 and TPR042 – Hybrid Information Feature.

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

#### 13.6.5.4 [\[Editor's note 13.6.6.5\]](#)HYBRID CHANGE BY LBA RANGE Subcommand

##### 13.6.5.4.1 [\[Editor's note 13.6.6.5.1\]](#)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 Editor's note 13.7.9.14 in TPR042).

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 13.6.5.4.2) 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.~~

[The values of the Max Priority Behavior bit \(see 13.6.5.4.2\), the Hybrid Priority field and the CB \(Cache Behavior\) bit shall control the movement of all of the selected logical sectors that are not already in the non-volatile caching medium into the non-volatile caching medium as defined in Editor's note Table 96+3\).](#)

Table 96+3 – Cache Behavior

<u>Hybrid Priority</u>	<u>Max Priority Behavior</u>	<u>Cache Behavior</u>	<u>Description</u>
<u>Maximum Priority</u>	<u>1</u>	<u>na</u>	<u>See [EDITOR'S NOTE 13.20.2.2.2 Max Priority Behavior set to one]</u>
<u>Maximum Priority</u>	<u>0</u>	<u>0</u>	<u>See [EDITOR'S NOTE 13.20.2.2.3 Max Priority Behavior is zero and Cache Behavior is zero]</u>
<u>Maximum Priority</u>	<u>0</u>	<u>1</u>	<u>See [EDITOR'S NOTE 13.20.2.2.4 Max Priority Behavior is zero and Cache Behavior is one]</u>
<u>Less than Maximum Priority but greater than zero</u>	<u>na</u>	<u>0</u>	<u>See [Editor's note 13.20.2.2.3 Hybrid Priority field is less than Maximum Hybrid Priority Level and is greater than zero]</u>
<u>Less than Maximum Priority but greater than zero</u>	<u>na</u>	<u>1</u>	<u>Device shall not copy any of the selected logical sectors that are not already in the non-volatile caching medium into the non-volatile caching medium if such movement causes the device to spin-up.</u>
<u>0</u>	<u>na</u>	<u>na</u>	<u>See [Editor's note 13.20.2.2.4 Hybrid priority is zero]</u>

If:

- a) the Hybrid Priority field (see Editor's note 13.20.6.5 in TPR042) is set to the Maximum Priority;
- b) the MaxPriority Behavior bit (see 13.6.5.4.2) 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			CB	3h			
Features(15:8)	Sector Count (7:0)							
Count(7:0)	TAG				Reserved			
Count(15:8)	Sector Count (15:8)							
LBA(7:0)	StartingLBA							
LBA(15:8)								
LBA(23:16)								
LBA(31:24)								
LBA(39:32)								
LBA(47:40)								
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**

**CB** [The CB \(Cache Behavior\) bit specifies the movement or non-movement of logical sectors into the non-volatile caching medium if the Hybrid Information field specifies a hybrid priority level other than the Max Priority. If the CB \(Cache Behavior\) bit is set to one, then the device shall not copy any of the selected logical sectors that are not already in the non-volatile caching medium into the non-volatile caching medium if such movement would cause spin-up of the device. If the CB \(Cache Behavior\) bit is cleared to zero and 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.](#)

**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 Editor’s note 13.20.6 in TPR042).

**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.

### 13.6.5.4.2 *[Editor's note 13.7.9.2.11]*Supported Options

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

Table 95+3 - Supported Options

Bit	Description
7: <del>4</del> 2	Reserved
<u>1</u>	<a href="#">Supports Cache Behavior</a>
0	Max Priority Behavior

[If the Supports Cache Behavior bit is set to one, then the device supports per command tagging of the HYBRID CHANGE BY LBA RANGE command to control the movement of logical sectors into the non-volatile caching medium if a command specifies a hybrid priority level other than the Max Priority \(see 13.6.5.4.1\).](#)

If the Max Priority Behavior bit is set to one (see Editor's note 13.20.5.5.2.2 in TPR042), 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 Editor's note 13.6.13.10 in TPR042); and
- d) support the HYBRID CHANGE BY LBA RANGE command (see 13.6.5.4).

If the Max Priority Behavior bit is cleared to zero (see Editor's note 13.20.5.5.2.2 in TPR042), then the device should insert logical sectors into the non-volatile caching medium ~~when~~if a command specifies the maximum hybrid priority level.

#### 13.20.5.5.2.2 [EDITOR'S NOTE 13.20.2.2.3]Max Priority Behavior is [cleared to zero and Cache Behavior is cleared to zero](#)

If:

- a) the device processes a read command or a write command, with the Hybrid Information Is Valid [bit](#) 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; ~~and~~ [and](#)
- c) [the Cache 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; [and](#)
- 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.



### **13.20.5.5.2.3 [EDITOR'S NOTE 13.20.2.2.4]Max Priority Behavior is cleared to zero and Cache Behavior is set to one**

#### **13.20.5.5.2.3.1 [EDITOR'S NOTE 13.20.2.2.4.1]Disk is spun down**

If:

- a) the device processes a read command or a write command, with the Hybrid Information Is Valid bit 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;
- c) the cache Behavior bit is set to one; and
- d) the disk is spun down,

then the device shall:

- a) not put the requested logical sectors in the non-volatile caching medium;
- b) not associate the logical sectors with the Hybrid Priority value specified by the command; and
- c) return command complete with no error.

#### **13.20.5.5.2.3.2 [EDITOR'S NOTE 13.20.2.2.4.2]Disk is spinning**

If:

- a) the device processes a read command or a write command, with the Hybrid Information Is Valid bit set to one and the Hybrid Priority value set to the Maximum Hybrid Priority Level;
- b) the Max Priority Behavior bit is cleared to zero;
- c) the Cache Behavior bit is set to one; and
- d) the disk is spinning,

then:

- a) the device shall behave as per section 13.20.5.5.2.2 [EDITOR'S NOTE 13.20.2.2.3 Max Priority Behavior is zero and Cache Behavior is zero]

### ***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***

#### **13.20.5.4.3.1 [EDITOR'S NOTE 13.20.2.2.3.1] Cache Behavior is cleared to 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, and the cache Behavior bit is cleared to zero, then:

- 1a) 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;
- 2b) 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

- 3e) the device should associate the logical sectors specified by the command with the Hybrid Priority value specified by the command; and
- 4) the device shall return command complete with no error.

#### **13.20.5.4.3.2 [EDITOR'S NOTE 13.20.2.2.3.2]Cache Behavior is set to one**

##### **13.20.5.4.3.2.1 [EDITOR'S NOTE 13.20.2.2.3.2.1]Disk is spun down**

If:

- a) the device processes a read command or a write command, with the Hybrid Information Is Valid bit set to one and the Hybrid Priority value set less than the Maximum Hybrid Priority Level;
- b) the Cache Behavior bit is set to one; and
- c) the disk is spun down,

then the device shall:

- a) not put the requested logical sectors in the non-volatile caching medium;
- b) not associate the logical sectors with the Hybrid Priority value specified by the command; and
- c) return command complete with no error.

##### **13.20.5.4.3.2.2 [EDITOR'S NOTE 13.20.2.2.3.2.2]Disk is spinning**

If:

- a) the device processes a read command or a write command, with the Hybrid Information Is Valid bit set to one and the Hybrid Priority value set less than the Maximum Hybrid Priority Level;
- b) the Cache Behavior bit is set to one; and
- c) the disk is spinning,

then:

- a) the device shall behave as if the Cache Behavior bit is cleared to zero (see **[EDITOR'S NOTE 13.20.2.2.3.1 Cache Behavior is cleared to zero]**).