Proposed **Serial ATA**

Draft **International Organization**

Version 12

November 26, 2012

TP_045v11_SATA31_TPR_D144

Title: Rebuild Assist

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.

Author Information

Author Name	Company	e-mail address
Jim Hatfield	Seagate Technology	james.c.hatfield@seagate.com

Workgroup Chair Information

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

Document History

Version	Date	Comments
0	07/06/2011	Initial draft
1	09/21/2011	Include feedback from July and September 2012 T10/CAP meetings 1. consider ALL resets 2. don't use FUA field of the read command to override rebuild assist mode 3. define a commanded way to disable the mode (not just power cycle) 4. restructured for more clarity
2	10/06/2011	1. Resolved some TBDs 2. Added support for errors writing to a disabled physical element 3. Renamed 'Next Accessible LBA' to 'Final LBA In Error' to be more consistent with reporting the location of the last error rather than the next expected good LBA
3	11/11/2011	1. (changes were not listed)
4	01/20/2012	1. Synchronized with T10/11-298r4 2. Specified handling of predicted write errors 3. Require NCQ Autosense to be supported 4. Removed modifications to WRITE FPDMA QUEUED command

		5. Simplified the requirements about error reporting
5	01/26/2012	1. Clarify RARC usage 2. More synching with T10/11-298r5 (minor overview wording changes) 3. On enabling the feature, the device 'may' (not 'shall') perform a diagnostic 4. Allow reporting of predicted errors in the Queued Error log for WRITE FPDMA QUEUED 5. Clarify how the host updates the Disabled Physical Elements field by logically OR'ing
6	02/10/2012	1. Removed material red-lined in previous reviews 2. Change copyright date from 2011 to 2012 3. Clarification about bit TBD3 in IDENTIFY DEVICE 4. Clarified requirement that adding bits to the Disabled Physical Elements field that are not supported shall result in command aborting the write to the log
7	03/14/2012	1. Synch with changes requested at March 2012 T10 plenary 2. Change: 'unrecoverable' to 'unrecovered' 3. Change 'corresponding to' to 'associated with' 4. Additional clarifications about behavior for unpredicted vs. predicted errors
8	05/16/2012	1. Synch with misc changes requested at May 2012 T10 plenary
9	06/18/2012	1. Misc editorial & style comments from the SATA Technical Editor (no technical changes)
10	10/05/2012	Changes from member review comments 1. change the log address from 14h to 15h 2. added description of changes to the GPL Log Directory contents 3. Queued Error log (10h): changed byte offset 9 from "LBA(39:24)" to "LBA(39:32)". This error was introduced in SATA Revision 3.0. 4. changed references from "ACS-2" to "ACS-3"
11	10/05/2012	Missed one more log address that had to be changed.

12	11/26/2012	1. Changed a ';' to a period at the end of a list 2. added a caption to the table of GPL log addresses

1 Introduction

This describes a new feature to be added to the next revision of the specification after SATA Revision 3.1.

2 Summary of the problem

A RAID rebuild operation is performed when some or all of a drive has failed in a RAID storage subsystem and a new drive is being initialized to replace the failed drive. The data for the new drive can be reconstructed by reading from the remaining drives in the RAID system and XORing the data to get the data that should be written to the replacement drive. This is time consuming but is the only way to recover data that is unrecovered.

This process can be performed faster if only part of the data on the failed drive is unreadable. The rebuild process is often hindered by the drive's normal recovery processes, however. If the drive takes several seconds to determine each LBA is unreadable and there are many unreadable Logical sectors grouped together, then this process can take longer than doing the rebuild from the other drives in the array. This proposal defines a way to shorten the long recovery process, so less time is wasted on unreadable Logical sectors, and provides a method for the drive to indicate quickly not only where the first LBA in error is, but also (if there are contiguous LBAs in error) where the last LBA in error is located.

Example of Rebuild Assist operations

This example presumes that Logical sectors are arranged in a serpentine format with 1000 blocks per track, and two heads (Figure 1).

- Track 0: Logical sectors 0 - 999 are mapped to head 0
- Track 1: Logical sectors 1000 - 1999 are mapped to head 1
- Track 2: Logical sectors 2000 - 2999 are mapped to head 0

Figure 1: LBA layout for Rebuild Assist example.

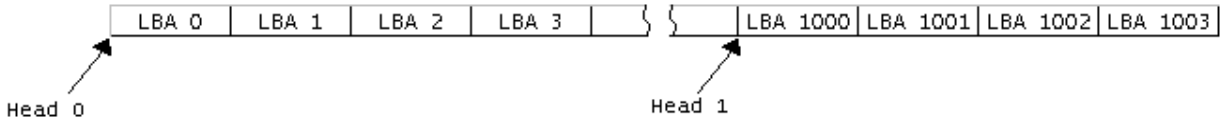

The host enters Rebuild Assist feature (WRITE LOG EXT command). Head 1 fails internal testing during the Rebuild Assist characterization process and is excluded. The Disabled Physical Element Mask becomes 0x00000002 (bit 1, representing head 1, is set to one).

The host issues a read command (Read FPDMA Queued) for 800 blocks starting at LBA 0. Since this range of LBA maps to head 0 (not excluded), all the data is returned; the command completes successfully.

The host then issues a read command for 800 blocks starting at LBA 800. Data for Logical sectors 800-999 is returned to the host, followed by an error response (Status 0x41 Error 0x24) which indicates that the remainder of the command could not be completed, since it requires access to an excluded head.

At this point, the error LBA information available to the host (Queued Error Log, Error, LBA) contains a value of 1000, which is the first LBA on an excluded head (head 1) that the command encountered. The Final LBA In Error (Queued Error Log, Final LBA In Error field) is 2000, the next sequential LBA on an included head (head 0). The host may continue reading starting at LBA 2000 (the Final LBA In Error).

3 Proposed changes

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

The following changes are proposed:

1. adding a new section to clause 13 to describe the Rebuild Assist feature;
2. modifications to IDENTIFY DEVICE command;
3. modifications to the READ FPDMA QUEUED command;
4. modifications to the Queued Error log;
5. adding a new log to section 13.7: Rebuild Assist log

[editors note: Insert this as a new subclause in clause 13]

13.xxx Rebuild Assist

13.xxx.1 Rebuild assist overview

The Rebuild Assist mode provides a method for a host controlling the rebuild process to determine which logical blocks on the failed device are unreadable without having to read every LBA to determine the unreadable logical blocks (i.e., the read command is terminated with an error and the failed LBA is reported in the sense data). The storage array controller then may reconstruct the failed logical blocks. The remaining logical blocks may be copied to the replacement device.

Enabling the Rebuild Assist feature:

a) may cause the device to initiate a self test to identify the scope of failures, if any;
b) modifies read command recovery behavior based on the setting of the RARC bit (see 13.6.2.2xxx); and
c) allows the NCQ Error log to indicate the location of multiple failing LBAs on both READ FPDMA QUEUED commands or WRITE FPDMA QUEUED commands.

If the device processes a power cycle, then the Rebuild Assist feature shall be disabled.
All other resets shall not affect the Rebuild Assist feature.

If the Rebuild Assist feature is supported, then the device shall also support the NCQ Autosense feature.

Self-test operations performed while Rebuild Assist mode is enabled may result in detection of failed physical elements.

A predicted unrecovered error is an unrecovered error that is the result of an attempt to access an LBA associated with a failed physical element.

An unpredicted unrecovered error is an unrecovered error that is the result of accessing an LBA that is not associated with a failed physical element.

13.xxx.2 Enabling Rebuild Assist feature

If the host writes to the Rebuild Assist log and sets the Rebuild Assist Enabled field to one, then:

a) the device may initiate a self test of the physical elements contained within the device and should disable any physical elements that are not functioning correctly;
b) the device shall initialize the Disabled Physical Elements from the results of the self-test;
c) the device shall minimize device-initiated background activities; and
d) the device shall enable the Rebuild Assist feature.

The host may verify that Rebuild Assist feature is enabled by reading the Rebuild Assist log, and then examining the data returned and verifying that the Rebuild Assist Enabled field is set to one.

13.xxx.3 Using the Rebuild Assist feature

13.xxx.3.1 Overview

If the Rebuild Assist feature is enabled, then the host should issue sequential READ FPDMA QUEUED commands to extract the available data from the device. If a READ FPDMA QUEUED command does not detect an unrecovered error, then the command should complete without error.

The Rebuild Assist feature allows reporting of an unrecovered read error or an unrecovered write error that is either predicted (i.e., a predicted unrecovered error) or unpredicted (i.e., an unpredicted unrecovered error).

If a device processes a READ FPDMA QUEUED command with the RARC bit set to one, then Rebuild Assist feature shall not affect processing of the READ FPDMA QUEUED command.

13.xxx.3.2 Unpredicted unrecovered read error

If the device processes a READ FPDMA QUEUED command with the RARC bit cleared to zero and detects an unpredicted unrecovered error, then the device:

a) performs limited read recovery that is vendor specific;
b) transfers the data for all recovered logical blocks, if any, from the starting LBA of the failed READ FPDMA QUEUED command up to the unrecovered logical block;
c) shall terminate the READ FPDMA QUEUED command with an error, with the following information recorded in the Queued Error log (see 13.7.3xxx):
    A) the Sense Key field shall be set to MEDIUM ERROR;
    B) the Additional Sense Code field and the Additional Sense Code Qualifier field shall be set to UNRECOVERED READ ERROR; and
    C) the LBA field shall be set to the LBA of the first unrecovered logical block; and
d) may use this failure in a vendor specific manner to predict other logical blocks that may be unrecovered.

NOTE: If the host receives sense data with sense key set to MEDIUM ERROR and additional sense code set to UNRECOVERED READ ERROR, then the host should issue the next read command with the starting LBA set to the contents of the Final LBA In Error field plus one.

13.xxx.3.3 Predicted unrecovered read error

If the device processes a READ FPDMA QUEUED command with the RARC bit cleared to zero and detects a predicted unrecovered error, then the device:

a) performs limited read recovery that is vendor specific;
b) transfers the data for all recovered logical blocks, if any, from the starting LBA of the failed READ FPDMA QUEUED command up to the first unrecovered logical block; and
c) shall terminate the READ FPDMA QUEUED command with an error, with the following information recorded in the Queued Error log (see 13.7.3xxx):
    A) the Sense Key field shall be set to ABORTED COMMAND;
    B) the Additional Sense Code field and the Additional Sense Code Qualifier field shall be set to MULTIPLE READ ERRORS;
    C) the LBA field shall be set to the LBA of the first unrecovered logical block; and
    D) the Final LBA In Error field shall be set to the LBA of the last predicted unrecovered logical block in a sequence of contiguous unrecovered logical blocks that started with the first LBA in error.

NOTE: If the host receives sense data with sense key set to ABORTED COMMAND and additional sense code set to MULTIPLE READ ERRORS, then the host should issue the next read command with the starting LBA set to the contents of the Final LBA In Error field plus one.

13.xxx.3.4 Unpredicted unrecovered write error

If the device encounters an unpredicted unrecovered error on a write command that is not the WRITE FPDMA QUEUED command, then the device shall terminate the command with an error.

If the device encounters an unpredicted unrecovered error on a WRITE FPDMA QUEUED command, then the device shall terminate the command with an error, with the following information recorded in the Queued Error log (see 13.7.3xxx):

a) the Sense Key field shall be set to MEDIUM ERROR;
b) the Additional Sense Code field and the Additional Sense Code Qualifier field shall be set to WRITE ERROR; and
c) the LBA field shall be set to the LBA of the first unrecovered logical block.

13.xxx.3.5 Predicted unrecovered write error

If the device encounters a predicted unrecovered error on a write command that is not the WRITE FPDMA QUEUED command, then the device shall terminate the command with an error.

If the device encounters a predicted unrecovered error on a WRITE FPDMA QUEUED, then the device shall terminate the command with an error, with the following information recorded in the Queued Error log (see 13.7.3xxx):

a) the Sense Key field shall be set to ABORTED COMMAND;
b) the Additional Sense Code field and the Additional Sense Code Qualifier field shall be set to MULTIPLE WRITE ERRORS;
c) the LBA field shall be set to the LBA of the first unrecovered logical block; and
d) the Final LBA In Error field shall be set to the LBA of the last predicted unrecovered logical block in a sequence of contiguous unrecovered logical blocks that started with the first LBA in error.

NOTE: If the host receives sense data with sense key set to ABORTED COMMAND and additional sense code set to MULTIPLE WRITE ERRORS, then the host should issue the next write command with the starting LBA set to the contents of the Final LBA In Error field plus one.

13.xxx.4 Disabling the Rebuild Assist feature

The Rebuild Assist feature shall be disabled if:

a) the device processes a power cycle; or
b) the device processes a command to write to the Rebuild Assist log (see 13.7.xxx) and clears the Rebuild Assist Enabled field to zero.

13.xxx.5 Testing the Rebuild Assist feature

The Rebuild Assist log (see 13.7.xxx) provides a method to test the host's rebuild process.

A device is put into a simulated failing condition by writing to the Rebuild Assist log with the ENABLED bit set to one and the Disabled Physical Elements field with one or more bits set to one. The host may write to the Rebuild Assist log more than once to simulate additional failing physical elements.

Each bit in the Disabled Physical Elements field represents a physical element that is associated with a group of LBAs that are treated as predicted unrecovered read errors and predicted unrecovered write errors. The correlation of bits in the Disabled Physical Elements field to LBAs in the device is vendor specific.

To end this test, disable the Rebuild Assist feature (see 13.x.4).

[editors note: modify the IDENTIFY DEVICE command as noted below]

13.2.1 IDENTIFY DEVICE

Table 1 – IDENTIFY DEVICE information

Word	O/M	F/V		
78	O	F	11	Rebuild Assist supported
79	O	V	11	Rebuild Assist enabled

13.2.1.18 Word 78: Serial ATA features supported

Bit 11, if set to one indicates that the device supports the Rebuild Assist feature (see 13.xxx). This bit shall only be set to one if the device supports NCQ as shown in bit 8 of Word 76. The host may determine if the Rebuild Assist feature is enabled or disabled by reading the Rebuild Assist log or by reading IDENTIFY DEVICE data word 79 bit 11.

13.2.1.19 Word 79: Serial ATA features enabled

Bit 11, if set to one indicates that the Rebuild Assist feature (see 13.xxx) is enabled. This bit shall only be set to one if the device supports the Rebuild Assist feature (i.e., IDENTIFY DEVICE data word 78 bit 11 is set to one) and the device supports NCQ (i.e., IDENTIFY DEVICE data word 76 bit 8 is set to one).

[editors note: modify the READ FPDMA QUEUED command as noted below]

13.6.2.2 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 [Editors note: Figure 224].

Register	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~~		
						Reserved		RARC
Count(15:8)	PRIO(1:0)		Reserved					
LBA(7:0)	LBA 7:0							
LBA(31:24)	LBA 31:24							
LBA(15:8)	LBA 15:8							
LBA(39:32)	LBA 39:32							
LBA(23:16	LBA 23:16							
LBA(47:40)	LBA 47:40							
ICC	ICC(7:0)							
Auxiliary(7:0)	Reserved							
Auxiliary (15:8)	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 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 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 QUEUE MANAGEMENT command (see [Editor's note 13.6.5]) and supports the Deadline Handling subcommand (see [Editor's note 13.6.5.1.3]). 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

- The time interval is fine-grained.
- Interval = 10 msec
- Time Limit = (ICC[6:0] + 1) * 10 msec
- Max Fine Time = 128 * 10 msec = 1.28 sec

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

- The time interval is coarse-grained
- Interval = 0.5 sec;
- Time Limit = (ICC[6:0] + 1) * 0.5 sec
- Max Coarse Time = 128 * 0.5 sec = 64 sec

FUA — When set to one forces the data to be retrieved from the storage media regardless of whether the storage device holds the requested information in its buffers or 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 media before being retrieved from the storage media as part of this operation. When cleared to zero the data may be retrieved either from the device's storage media or from buffers/cache that the device may include.

RARC — If the Rebuild Assist feature is not supported, then the RARC bit shall be ignored.

If the Rebuild Assist feature is supported and is disabled, then the RARC bit shall be ignored.

If the Rebuild Assist feature is supported and enabled, then the RARC bit specifies that read operations shall be processed as defined in 13.xxx.

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.

13.6.2.3 Success Outputs

[editors note: no change to this subclause]

13.6.2.4 Error Outputs

[editors note: no change to this subclause]

13.6.2.5 Queue abort

[editors note: no change to this subclause]

[editors note: modify subclause 13.7.1 as noted below]

13.7.1 Log Address Definitions

The log addresses assigned for Serial ATA are defined in [editors note: Table 93].

Table ~~94~~ 93 - Log Addresses for Serial ATA

Log Address	Description
00h - 0Fh	As defined in the ~~ATA8-ACS~~ ACS-3 standard
10h	NCQ Queued Error log
11h	Phy Event Counters log
12h	NCQ Queue Management log
13h	NCQ Send and Receive log
14h	Reserved
~~14h~~15h	Rebuild Assist log
16~~5~~h~~14h~~ - 17h	Reserved
18h - FFh	As defined in the ~~ATA8-ACS~~ ACS-3 standard

[editors note: modify subclause 13.7.2 as noted below]

13.7.2 General Purpose Log Directory (00h)

Devices supporting the Queued Error Log (see 13.7) reflect this support in the General Purpose Log Directory 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 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 Rebuild Assist log reflect this support in the General Purpose Log Directory (00h) by having the value 1 at offset 2Ah and the value 0 at offset 2Bh of that log to indicate existence of a log at address 15h of 1 page in length.

Table 94 - General Purpose Log directory values for Serial ATA

Byte	Log	Value
000h..01Fh		As defined in the ACS-3 standard
020h	10h	1 if Native Command Queuing is supported, 0 if Native Command Queuing is not supported
021h	10h	0
022h	11h	1 if Phy Event Counters are supported 0 if Phy Event Counters are not supported
023h	11h	0
024h	12h	1 if NCQ Queue Management is supported 0 if NCQ Queue Management is not supported
025h	12h	0
026h	13h	1 if NCQ Send and Receive log is supported 0 if NCQ Send and Receive log is not supported
027h	13h	0
028h		Reserved
029h		Reserved
02Ah	15h	1 if Rebuild Assist log is supported 0 if Rebuild Assist log is not supported
02Bh	15h	0
02Ch..2Fh		Reserved
030h..1FFh		As defined in the ACS-3 standard

[editors note: modify subclause 13.7.3 Queued Error Log as noted below]

13.7.3 Queued Error Log (10h)

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

The General Purpose Logging (GPL) feature set is defined in the ~~ACS-2~~ACS-3 standard.

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

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

Reading the Queued Error Log (10h) has the additional side effect defined in section [Editor's note 13.6.2] of aborting any outstanding queued commands and returns a device that has halted due to a queued command error to a state where it has no commands outstanding and is again ready to accept commands (e.g., after completion of a command to read the log the device returns to state DI0:Device_idle state as defined in section [Editor's note 11.2]). The Queued Error Log contains extended command error information.

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

Devices supporting the native queued capability shall support the Queued Error Log. The Queued Error Log is one page in length and is defined in [Editor's note: Figure 252].

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

a) the Sense Key field;
b) the Additional Sense Code field; and
c) the Additional Sense Code Qualifier field,

shall be set to values defined in SPC-4.

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

a) the Sense Key field shall be cleared to zero;
b) the Additional Sense Code field shall be cleared to zero; and
c) the Additional Sense Code Qualifier field shall be cleared to zero.

Byte	7	6	5	4	3	2	1	0
0	NQ	UNL	R	TAG				
1	Reserved							
2	Status							
3	Error							
4	LBA(7:0)							
5	LBA(15:8)							
6	LBA(23:16)							
7	Device							
8	LBA(31:24)							
9	LBA(39:~~24~~32)							
10	LBA(47:40)							
11	Reserved							
12	Count(7:0)							
13	Count(15:8)							
14	Sense Key							
15	Additional Sense Code							
16	Additional Sense Code Qualifier							
17	~~Reserved~~ Final LBA In Error(7:0)							
18	~~Reserved~~ Final LBA In Error(15:8)							
19	~~Reserved~~ Final LBA In Error(23:16)							
20	~~Reserved~~ Final LBA In Error(31:24)							
21	~~Reserved~~ Final LBA In Error(39:32)							
22	~~Reserved~~ Final LBA In Error(47:40)							
23~~20~~ … 255	Reserved							
256 … 510	Vendor Specific							
511	Data Structure Checksum							

Figure 252 – Queued Error Log data structure definition

TAG If the NQ bit is cleared to zero, the TAG field contains the TAG corresponding to the queued command that failed.

UNL If set to one indicates that the error condition was a result of receiving an IDLE IMMEDIATE command with the Unload Feature specified. If cleared to zero, the reason for the error was not due to reception of an IDLE IMMEDIATE command

with the Unload Feature specified. If the last command received was an Unload Immediate, the device shall not load the heads to the media when reading the Queued Error Log.

If set to one, the NQ bit shall also be set to one to indicate the failure was due to reception of a non-queued command. When set to one, the value of the Status, Error, and LBA(7:0) fields (bytes 3-5) in the log shall be set as follows:

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

Error: ABRT bit shall be set to one

LBA(7:0): Shall be set to C4h if the unload is being executed or has completed successfully. Shall be set to 4Ch if the unload was not accepted or has failed.

NQ If set to one indicates that the error condition was a result of a non-queued command having been issued and that the TAG field is therefore not valid. If cleared to zero indicates that the TAG field is valid and that the error condition applies to a queued command.

BYTE1-19 An image of a Register Device to Host FIS is embedded in the data structure. The fields correspond to the Shadow Register Block Registers and are encoded with error information consistent with the READ DMA QUEUED EXT or WRITE DMA QUEUED EXT command defined in the ATA8-ACS standard.

ERROR The value corresponding to the ATA ERROR register value for the command that failed. The command-specific error condition of invalid tag value shall be handled as an invalid command parameter and shall be reported as such (i.e. ABRT bit set to one in the error register and all other bits cleared to zero).

Note that the value returned in the ERROR field of the data structure is separate from the value returned in the Error shadow register when the initial error condition is signaled. The Error shadow register value is used for the purpose of signaling a queued command error, while the value in the ERROR field of the data structure provides specific information about the error condition that the specific queued command encountered.

~~Vendor Specific~~

~~Allocated for vendor specific use.~~

Sense Key

See SPC-4.

Additional Sense Code

See SPC-4

Additional Sense Code Qualifier

See SPC-4

Final LBA In Error

If:

a) the command in error is READ FPDMA QUEUED or WRITE FPDMA QUEUED;

b) the Sense Key is ABORTED COMMAND; and

c) the Additional Sense Code/Additional Sense Code Qualifier is MULTIPLE READ ERRORS or MULTIPLE WRITE ERRORS,

then the Final LBA In Error field shall contain the LBA of the last logical block in a sequence of contiguous unrecovered logical blocks. Otherwise, the Final LBA In Error field shall be cleared to zeroes.

Vendor Specific

Allocated for vendor specific use.

Data Structure Checksum

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

Reserved/R

All reserved fields shall be cleared to zero.

[editors note: insert this new subclause into subclause 13.7]

13.7.xxx Rebuild Assist log (15h)

If the device supports the Rebuild Assist feature (i.e., IDENTIFY DEVICE data word 78 bit 11 is set to one), then the Rebuild Assist log shall be supported.

The Rebuild Assist log shall be accessed using the GPL feature set commands (see ATA8-ACS).

If the Rebuild Assist log is not supported and the host:

a) reads the Rebuild Assist log; or

b) writes the Rebuild Assist log,

then the device shall return command aborted.

The Rebuild Assist log provides information about the Rebuild Assist feature (see 13.xxx).

If the host writes to the Rebuild Assist log and the device supports the Rebuild Assist feature and the Rebuild Assist Enabled field is cleared to zero, then the device shall:

1) disable the Rebuild Assist feature;
   a) clear the Rebuild Assist Enabled field to zero;
   b) set the Length of the Disabled Physical Element Mask field to a non-zero value;
   c) set the Disabled Physical Element Mask field to a vendor specific value; and
   d) clear the Disabled Physical Elements field to zero;
2) ignore all other data from the host; and
3) return command completion with no error.

If the host writes to the Rebuild Assist log and the device supports the Rebuild Assist feature and the Rebuild Assist Enabled field is set to one, then:

1) if:
   a) the device is unable to enable the Rebuild Assist feature;
   b) the host attempts to set any bits to one in the Disabled Physical Elements field that are cleared to zero in the Disabled Physical Element Mask field; or
   c) the host attempts to set all bits to one in the Disabled Physical Elements field that are set to one in the Disabled Physical Element Mask field (i.e., attempt to disable all physical elements),

   then the device shall return command aborted;

2) the device shall enable the Rebuild Assist feature (see 13.xxx.2);
3) if the device successfully enabled the Rebuild Assist feature, then the device shall logically OR the Disabled Physical Elements field with any prior Disabled Physical Elements field that the device was using (e.g., the host may add bits but shall not clear bits in the field) and save the new value of the Disabled Physical Elements field; and
4) the device shall set IDENTIFY DEVICE word 79 bit11 to one (i.e., Rebuild Assist feature enabled).

If the host reads from the Rebuild Assist log and Rebuild Assist feature is supported, then:

a) if the Rebuild Assist feature is disabled, then the device shall return the current values for all fields; and

b) if the Rebuild Assist feature is enabled, then the device:

A) shall set the Rebuild Assist Enabled field to one;

B) may set additional bits in the Disabled Physical Elements field; and

C) shall not clear any Disabled Physical Elements bits that were previously set by the host.

Byte	Description
0	Flag Bits Bits 7:1 Reserved 0 Rebuild Assist Enabled
1..6	Reserved
7	Physical Element Length (N)
8	(MSB)
	Disabled Physical Element Mask
7+N	(LSB)
8+N	(MSB)
	Disabled Physical Elements
7+(2×N)	(LSB)
8+(2×N) .. 511	Reserved

Figure 1 - Rebuild Assist Log

Rebuild Assist Enabled

Table 2 describes the use of the Rebuild Assist Enabled field.

Table 2 - Rebuild Assist Enabled field

Operation	Rebuild Assist Enabled	Description
read log	1	the Rebuild Assist feature is enabled
read log	0	the Rebuild Assist feature is disabled
write log	1	request to enable the Rebuild Assist feature
write log	0	request to disable the Rebuild Assist feature

Physical Element Length

The Physical Element Length field indicates the number of bytes in the Disabled Physical Element Mask field and the number of bytes in the Disabled Physical Elements field.

The device shall ignore any attempt by the host to change the value of this field when writing to the Rebuild Assist log.

Disabled Physical Element Mask

The Disabled Physical Element Mask field indicates which bits in the Disabled Physical Elements field are supported.

The device shall ignore any attempt by the host to change the value of this field when writing to the Rebuild Assist log.

Disabled Physical Elements

The Disabled Physical Elements field specifies if physical elements shall be disabled. Each bit that is set to one in the Disabled Physical Elements field specifies that LBAs associated with this physical element shall respond to read commands and write commands as if the associated LBAs have predicted errors (see 13.xxx). Each bit that is set to zero in the Disabled Physical Elements field specifies that LBAs associated with this physical element shall respond to read commands and write commands as if the associated LBAs do not have predicted errors.

