

TRIM – DRAT / RZAT clarifications for ATA8-ACS2

~~December~~ February 2314, 201009

Revision 2~~1~~

Technical Editor:

Curtis E. Stevens
20511 Lake Forest Dr.
Lake Forest, Ca. 92630
949-672-7933
Curtis.Stevens@wdc.com

Author:

Fred Knight
7301 Kit Creek Rd
PO Box 13917
Research Triangle Park, NC 27709
919-476-5362

Email: knight@netapp.com

Document Status

Revision History (part of)		
Rev	Date	Description
1	December 14, 2009	1) Initial Revision
<u>2</u>	<u>February 23, 2010</u>	<u>1) Incorporate comments from Feb Plenary meeting</u>

1 Introduction

Clarification of interactions of TRIM with DRAT and RZAT was requested during letter ballot review.

2 Scope

This proposal introduces provides that clarification.

3 Overview

Clarify the interactions of TRIM with DRAT and RZAT.

TRIM bit – Identify Device word 169 bit 0

DRAT bit – Identify Device word 69 bit 14

RZAT bit – Identify Device word 69 bit 5

4 Changes to ACS

7.18.7.30 Word 69: Additional Supported

Word 69 shall indicate features/command sets supported. If a defined bit is cleared to zero, the indicated features/command set is not supported. These features/command sets are enabled and there is no disable mechanism.

If bit 15 of word 69 is set to one, then the device supports the CFast specification.

If word 169 bit 0 is set to one and word 69 bit 14 is cleared to zero, then the Trim function of the DATA SET MANAGEMENT command (see 7.10.3.2) supports indeterminate read after trim behavior. If word 169 bit 0 is set to one and word 69 bit 14 is set to one, the Trim function of the DATA SET MANAGEMENT command supports determinate read after trim behavior. If word 169 bit 0 is cleared to zero, then word 69 bit 14 is reserved.

If bit 13 of word 69 is set to one, then the device supports the SET FEATURES command with the Long Physical Sector Alignment Error Reporting Control subcommand (see 7.50.20).

If bit 12 of word 69 is set to one, then the DEVICE CONFIGURATION IDENTIFY DMA and DEVICE CONFIGURATION SET DMA commands are supported.

If bit 11 of word 69 is set to one, then the READ BUFFER DMA command is supported.

If bit 10 of word 69 is set to one, then the WRITE BUFFER DMA command is supported.

If bit 9 of word 69 is set to one, then the SET MAX SET PASSWORD DMA and SET MAX UNLOCK DMA commands are supported.

If bit 8 of word 69 is set to one, then the DOWNLOAD MICROCODE DMA command is supported.

Bit 7 is reserved for IEEE-1667

Bit 6 of word 69 shall be cleared to zero if the following commands are supported:

- a) FLUSH CACHE;
- b) READ DMA;

- c) READ MULTIPLE;
- d) READ SECTOR(S);
- e) READ VERIFY SECTOR(S);
- f) SET MULTIPLE MODE;
- g) WRITE DMA;
- h) WRITE MULTIPLE; and
- i) WRITE SECTOR(S).

Bit 6 of word 69 shall be set to one if any of the following commands are not supported:

- a) FLUSH CACHE;
- b) READ DMA;
- c) READ MULTIPLE;
- d) READ SECTOR(S);
- e) READ VERIFY SECTOR(S);
- f) SET MULTIPLE MODE;
- g) WRITE DMA;
- h) WRITE MULTIPLE; or
- i) WRITE SECTOR(S).

If word 469 bit 14 is set to one and word 69 bit 5 is set to one, then a read operation after a Trim operation returns data from trimmed LBAs as all words cleared to zero. If word 469 bit 14 is set to one and word 69 bit 5 is cleared to zero, then a read operation after a Trim operation may have words set to any value. If word 469 bit 14 is cleared to zero, then word 69 bit 5 is reserved.

Bits (4:0) of word 69 are reserved.

7.10 DATA SET MANAGEMENT – 06h, DMA

7.10.1 Feature Set

This 48-bit command is optional for ATA devices.

7.10.2 Description

The DATA SET MANAGEMENT command provides information for device optimization (e.g., file system information).

7.10.3 Inputs

7.10.3.1 Overview

Name	Description
Feature	<p>Bit Description</p> <p>15:1 Reserved.</p> <p>0 Trim - See 7.10.3.2.</p>
Count	Number of 512-byte blocks of LBA Range Entries to be transferred. The value of zero is reserved.
LBA	Reserved
Device	<p>Bit Description</p> <p>7 Obsolete</p> <p>6 Shall be set to one</p> <p>5 Obsolete</p> <p>4 Transport Dependent - See 6.2.12</p> <p>3:0 Reserved</p>
Command	7:0 06h

7.10.3.2 Trim

If the Trim bit is set to one, then a trim operation is being requested on the data in the logical block ranges addressed by the DATA SET MANAGEMENT command's output data. ~~shall.~~ Until a subsequent write command to a logical block that has been trimmed successfully completes, if the data from that logical block address is read, then the data is returned as described in table X.

- a) ~~if IDENTIFY DEVICE data word 69 bit 14 is set to one, then:~~
 - A) ~~after a trimmed LBA has been read (e.g., a read command has been processed), the data in that logical block becomes determinate (i.e., all read commands to a logical block shall return the same data until a subsequent write command to that logical block successfully completes);~~
 - B) ~~if IDENTIFY DEVICE data word 69 bit 5 is set to one, then the data returned by a read command shall have all words cleared to zero; and~~
 - C) ~~if IDENTIFY DEVICE data word 69 bit 5 is cleared to zero, then the data returned by a read command may have words set to any value;~~

~~or~~

- b) ~~if IDENTIFY DEVICE data word 69 bit 14 is cleared to zero, then the data read is indeterminate.~~

Table X shows the interactions of IDENTIFY DEVICE data word 169 bit 0, IDENTIFY DEVICE data word 69 bit 14, and IDENTIFY DEVICE data word 69 bit 5.

Table X – TRIM related interactions

<u>Word 169 bit 0</u>	<u>Word 69 bit 14</u>	<u>Word 69 bit 5</u>	<u>Description</u>
<u>0</u>	<u>Reserved</u>	<u>Reserved</u>	<u>The Trim function of the DATA SET MANAGEMENT command (see 7.10.3.2) is not supported. The data is unaffected by this command.</u>
<u>1</u>	<u>0</u>	<u>Reserved</u>	<u>The Trim function of the DATA SET MANAGEMENT command (see 7.10.3.2) shall cause indeterminate read after trim behavior (1).</u>
<u>1</u>	<u>1</u>	<u>0</u>	<u>The Trim function of the DATA SET MANAGEMENT command (see</u>

			<u>7.10.3.2) shall cause deterministic read after trim behavior (2) with words set to any value.</u>
<u>1</u>	<u>1</u>	<u>1</u>	<u>The Trim function of the DATA SET MANAGEMENT command (see 7.10.3.2) shall cause deterministic read after trim behavior (2) with words set to zero.</u>
<u>Footnote 1: Indeterminate read after trim behavior: each read command to the logical block may return different data.</u>			
<u>Footnote 2: Determinate read after trim behavior: after a read command has completed processing, the data in that logical block becomes determinate (i.e., all read commands to the logical block shall return the same data).</u>			

The data read from an LBA that has been trimmed shall not be retrieved from data that was previously received from an application client addressed to any other LBA.

Once a trimmed LBA has been written (e.g., a write command or a SECURITY ERASE UNIT command), the data in that logical block becomes determinate (i.e., the logical block contains the written data).

Trim shall not add or remove LBAs from the NV Cache Pinned Set.

7.10.4 Normal Outputs

See table 111.

7.10.5 Error Outputs

If trim is set to one and:

- a) the device detects an invalid LBA Range entry; or
- b) Count is greater than IDENTIFY DEVICE data word 105

then the device shall return command aborted.

A device may Trim one or more LBA Range Entries before it returns command aborted. See table 137.

7.10.6 Output From the Host to the Device Data Structure

DATA SET MANAGEMENT Request Data is a list of one or more LBA Range Entries (see 4.18.3.2). If the Range Length of an LBA Range Entry is zero then the LBA Range Entry is not valid. If Trim is set to one, then LBA Range Entries may overlap and are not required to be sorted. See table 44.