

Maximum LBA Range Supported Proposal ATA8-ACS2

February 17, December 2, 2009

Revision 41

Technical Editor:

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

Author:

Rob Strong
1900 Prairie City Road
Folsom, CA 95630
916-356-0616

Email: robert.strong@intel.com

Document Status

Revision History (part of)		
Rev	Date	Description
0	December 1, 2008	1) Initial Revision
1	December 2, 2008	2) Changed Count definition to specify the number of logical blocks instead of the number of LBA Range Entries.
<u>2</u>	<u>February 2, 2009</u>	<u>3) Updated based on feedback from December plenary. Added language for overlapped and non-sorted support. Added table to describe Trim Request Data</u>
<u>3</u>	<u>February 17, 2009</u>	<u>4) Updated with comments per February plenary (first review)</u>
<u>4</u>	<u>February 17, 2009</u>	<u>5) Updated revision 7.10.5 Error Outputs</u>

1 Introduction

~~ATA/ATAPI Command Set - 2 (ACS-2) The e07154r6 proposal defines the -added the-~~ DATA SET MANAGEMENT ~~06h, DMA~~ command and ~~associated the-~~ ~~Trim~~ ~~RIM~~ function.

The DATA SET MANAGEMENT command defines a Count field that allows the host to specify the number of logical blocks of LBA Range Entry to be transferred, where the minimum number of logical blocks that can be specified is one and the maximum is 65,536 (Count == 0). The current DATA SET MANAGEMENT command provides no mechanism to limit this ~~to a~~ value ~~to one~~ that is appropriate to the specific implementation.

For some device implementations this inability to limit the number of host provided LBA Range Entries could be problematic. As an example, some low cost implementations may not have enough local DRAM to store a 32MB set of LBA Range Entries.

2 Scope

This proposal introduces a mechanism for a device to indicate the maximum number of LBA Range Entries that the host may specify for each invocation of the DATA SET MANAGEMENT command.

3 Overview

This proposal introduces a WORD in the IDENTIFY DEVICE data to indicate the maximum number of LBA Range Entries that may be specified by the DATA SET MANAGEMENT command.

4 Changes to ACS-2

4.1 Definition of Terms

~~4.1.1 Invalid LBA: A LBA that is greater than the value reported in IDENTIFY DEVICE data words 100..103.~~

~~4.1.2 Invalid LBA Range: A range of LBAs that contains one or more invalid LBAs.~~

4.14.2 Changes to IDENTIFY DEVICE command

~~Table 29~~ ~~Table 30~~ — IDENTIFY DEVICE data changes:

ACS-2 IDENTIFY DEVICE ~~Words~~ ~~[TBD1]..[TBD2]~~ is used by the device during device enumeration to indicate the maximum number of LBA Ranges Entries supported. The chosen IDENTIFY DEVICE word shall be augmented with the following description:

Word	O M	S P	F V	Description
[TBD 1]..[TBD 2]	O	B	F	Maximum number of LBA Range Entries (See 4.16.3.2*) per DATA SET MANAGEMENT command (DWord).:

~~(*) LBA Range Entry as defined in ATA8-ASC (4.16.3.2 LBA Range Entry)~~

Add the following sub-clause:

7.16.7.XX Words ~~[TBD1]..[TBD2]~~: Maximum number of LBA Range Entries per DATA SET MANAGEMENT command

~~Words~~ ~~Word~~ ~~[TBD1]..[TBD2]~~ contains the maximum number of LBA Range Entries per DATA SET MANAGEMENT command that the ATA device shall accept. A value of ~~0000~~ 0000h indicates that the device ~~is able to can~~ process ~~4 194 304 a maximum of 65,536~~ LBA Range Entries per DATA SET MANAGEMENT command.

If word 169 bits ~~0~~[TBD3] is set to zero, the Trim function of the DATA SET MANAGEMENT _command is not supported and ~~therefore words~~word [TBD1]..~~[TBD2] are is~~ undefined.

DCO – No changes. The Maximum number of LBA Range Entries per DATA SET MANAGEMENT command is not changeable.

4.24.3 Changes to DATA SET MANAGEMENT – 06h, DMA command (~~per e07154r6~~)

The purpose of DATA SET MANAGEMENT command is for host to provide file system information for device optimization.

7.10 DATA SET MANAGEMENT - 06h, DMA

7.10.1 Feature Set

This 48-bit command is optional for devices that implement the General feature set.

7.10.2 Description

The purpose of the DATA SET MANAGEMENT command is for host to provide file system information for device optimization.

7.10.3 Inputs

7.10.3.1 Overview

~~ATA Command Format for DATA SET MANAGEMENT Command~~

Name	Description												
Feature	<table border="1"> <thead> <tr> <th>Bit</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>15:1</td> <td>Reserved</td> </tr> <tr> <td>0</td> <td>Trim – See 7.10.3.2</td> </tr> </tbody> </table>	Bit	Description	15:1	Reserved	0	Trim – See 7.10.3.2						
Bit	Description												
15:1	Reserved												
0	Trim – See 7.10.3.2												
Count	Number of 512-byte logical blocks of LBA Range Entries to be transferred, 0000h indicates that 65,536 512-byte logical blocks are to be transferred. This value is bound by word [TBD1] of the IDENTIFY DEVICE command data.												
LBA	Reserved												
Device	<table border="1"> <thead> <tr> <th>Bit</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>7</td> <td>Obsolete</td> </tr> <tr> <td>6</td> <td>Shall be set to one</td> </tr> <tr> <td>5</td> <td>Obsolete</td> </tr> <tr> <td>4</td> <td>Transport Dependent – See 6.2.12</td> </tr> <tr> <td>3:0</td> <td>Reserved</td> </tr> </tbody> </table>	Bit	Description	7	Obsolete	6	Shall be set to one	5	Obsolete	4	Transport Dependent – See 6.2.12	3:0	Reserved
Bit	Description												
7	Obsolete												
6	Shall be set to one												
5	Obsolete												
4	Transport Dependent – See 6.2.12												
3:0	Reserved												
Command	7:0 06h												

Word	Name	Description	
00h	Feature	Bit	Description
		15:1	Reserved
		0	Trim
01h	Count	Number of logical blocks of LBA Range Entries to be transferred, 0000h indicates that 65,536 logical blocks are to be transferred. This value is bound by word [TBD1] of the IDENTIFY DEVICE command data.	
02h 04h	LBA	Bit	Description
		47:0	Reserved
05h	Command	XXh	

Trim

~~When trim is set to one the data in the logical blocks specified by the DATA SET MANAGEMENT command's output data becomes indeterminate.~~

~~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.~~

7.10.3.2 Trim

When Trim is set to one the data in the logical blocks specified by the DATA SET MANAGEMENT command's output data then:

- a) if word 69 bit 14 of the IDENTIFY DEVICE data is set to one, then once a trimmed LBA has been read (e.g., a read command), 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); or
- b) if word 69 bit 14 of the IDENTIFY DEVICE data is cleared to zero, then the data read is indeterminate.

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 112.

~~**4.3** Normal Outputs~~
~~No changes.~~

~~**4.4** Error Outputs~~

~~**7.10.5 Error Outputs**~~

~~If trim is set to one and:~~

- ~~a. the device detects an invalid LBA Range entry;~~
- ~~b. Count is zero and (IDENTIFY DEVICE data words [tbd1]..[tbd2] is greater than zero and less than 4 194 241); or~~

c. Count is greater than zero and ((Count-1) * 64) is greater than IDENTIFY DEVICE data words [TBD1]..[TBD2], then the device shall return command aborted.

~~See table 138. When Trim is set to one and the device detects an invalid LBA Range Entry, then the device shall return command aborted. A device may successfully complete Trim RIMing one or more LBA Range Entries before it returns command aborted. command aborts due detection of an invalid LBA Range Entry value. See table 125.~~

~~When Trim is set to one, a device shall return command aborted if the host provided Count value is greater than that specified through IDENTIFY DEVICE word [TBD1]. See table 125.~~

7.10.6 Output From the Host to the Device Data Structure

~~See 7.22.3.6. Trim Request Data is a list of one or more individual LBA ranges. Each entry in the TRIM Request Data is called a LBA Range Entries (see 4.17.3.2). If the range length is zero then the LBA Range Entry is not valid. The range entries may overlap and are not required to be sorted. may be non-sorted. See Table 44 for the Trim Request Data structure.~~

Table TBD — TRIM Request Data

<u>Byte</u>	<u>Word</u>	<u>Description</u>
<u>0..7</u>	<u>QWORD</u>	<u>Entry #0</u> <u>63:48 Range Length</u> <u>47:0 LBA Value</u>
<u>8..15</u>	<u>QWORD</u>	<u>Entry #1</u> <u>63:48 Range Length</u> <u>47:0 LBA Value</u>
<u>...</u>	<u>...</u>	<u>...</u>
<u>496..511</u>	<u>QWORD</u>	<u>Entry #63</u> <u>63:48 Range Length</u> <u>47:0 LBA Value</u>

4.4

~~A device shall return command aborted if the device has received a DATASET MANAGEMENT command with a host provided Count value that is greater than that specified through IDENTIFY DEVICE word [TBD1]. See Table 123, Generic Abort Error.~~