

Non Volatile Cache Command Proposal for ATA8-ACS

~~Feb. Nov. 44~~th, 20065
Revision ~~65~~

Author: Nathan Obr
Frank Shu
Microsoft Corporation
One Microsoft Way
Redmond, WA. 98052-6399
USA

Phone: (425) 707-1779
E-Mail: frankshu@microsoft.com

Table of Contents

1	INTRODUCTION	<u>42</u>
2	DESCRIPTION OF NV CACHE	<u>42</u>
3	DESCRIPTION OF NV CACHE POWER MANAGEMENT	<u>42</u>
4	NV CACHE COMMAND PROPOSAL.....	<u>42</u>
4.1	PROPOSED SOLUTION REQUIREMENTS.....	<u>52</u>
4.1.1	<i>NV Cache Command Requirements</i>	<u>52</u>
4.1.2	<i>NV Cache Power Management Requirements.....</i>	<u>52</u>
4.2	PROPOSED DEFINITION OF TERMS FOR SOLUTION.....	<u>52</u>
4.3	SOLUTION OVERVIEW: SCENARIOS AND ILLUSTRATIONS	<u>62</u>
4.3.1	<i>NV Cache Sets</i>	<u>62</u>
4.3.2	<i>Pinning</i>	<u>62</u>
4.3.3	<i>Cache Management.....</i>	<u>72</u>
4.3.4	<i>Behavior over Power Cycle</i>	<u>82</u>
4.3.5	<i>Preparing to Pin a Large NV Cache Set Data</i>	<u>82</u>
5	PROPOSED CHANGES TO ATA8-ACS	<u>82</u>
5.1	IDENTIFY DEVICE	<u>82</u>
5.2	DEVICE CONFIGURATION OVERLAY FOR NV CACHE FEATURE SET	<u>102</u>
5.3	SECURITY FEATURE SET.....	<u>102</u>
5.4	LONG LOGICAL SECTOR FEATURE SET FOR NON-PACKET DEVICES.....	<u>112</u>
6	PROPOSED NEW NV CACHE POWER MANAGEMENT FOR ATA8-ACS	<u>112</u>
6.1	NV CACHE POWER MANAGEMENT FEATURE SET.....	<u>112</u>
6.1.1	<i>Set NV Cache Power Mode – B6h/00h, Non-data</i>	<u>122</u>
6.1.2	<i>RETURN From NV Cache Power Mode – B6h/01h, Non-data</i>	<u>132</u>
6.2	PROPOSED POWER STATE CHANGES	<u>132</u>
6.3	CHECK POWER MODE CHANGES.....	<u>142</u>
7	PROPOSED NEW NV CACHE COMMANDS FOR ATA8-ACS	<u>142</u>
7.1.1	<i>Add LBA(s) to NV Cache Pinned Set – B6h/10h, DMA.....</i>	<u>152</u>
7.1.2	<i>Remove LBA(s) From NV Cache Pinned Set – B6h/11h, DMA</i>	<u>172</u>
7.1.3	<i>Query NV Cache Pinned Set – B6h/12h, DMA</i>	<u>192</u>
7.1.4	<i>Query NV Cache Misses – B6h/13h, DMA.....</i>	<u>202</u>
7.1.5	<i>Flush NV Cache – B6h/14h, Non-data</i>	<u>212</u>

Revision History		
Date	Revision	Description

Revision History		
Date	Revision	Description
2005-02-07	0	Initial Draft
2005-04-01	1	Added proposed NV Cache commands and mode
2005-05-11	2	Incorporated April T13 feedback <ul style="list-style-type: none">• Removed NV Cache Power Mode• Reformed Resize command into Flush Command• Added scenarios and illustrations
2005-06-20	3	Updated formatting from ATA7 to ATA8-ACS
2005-10-09	4	Incorporating the feedback from June 05 T13 meeting
2005-11-17	5	Incorporating the feedback from Oct. 05 T13 meeting

1 Introduction

This proposal discusses the possibility of a Non Volatile (NV) Cache on an ATA device. Such a device would have the advantages of faster random access to the sectors specified by host , faster system boot times and improved power management due to its ability to satisfy reads and writes while the drive rotational media is spun down. This new hardware feature will also enable extended product life and improved shock resistance for the device.

In order for the device to take full advantage of an NV cache, the host would need to provide some configuration information about which sectors must be cached during different scenarios. This document provides a common understanding of the new concepts an NV cache introduces, provides a common language to discuss NV Cache, and proposes a new NV Cache command set to take advantage of such a device.

2 Description of NV Cache

The NV Cache effectively creates a tier of permanent storage between the host adaptor and the rotational media of the ATA device. This NV Cache can be considered to be behind a device's buffer but could also be implemented as complimentary to the device's buffer or possibly in a bridge before the device's buffer.

The addition of NV Cache doesn't extend the capacity of the storage device; instead it only duplicates a varying selection of the device capacity. The intent of this proposal is to create a method for managing the contents of the NV Cache. The addition of this new feature would never permit a device to satisfy a read of an LBA with anything but the value that was last written to that LBA.

Additionally this proposal doesn't extend beyond the ATA command set and provides no requirements about what internal mechanism is used to implement an NV Cache. It is conceivable that parts of the nonvolatile memory incorporating the logical NV Cache could be used by the device for storing its own operating information and that the device never exposes that capacity as being part of the NV Cache to the host.

3 Description of NV Cache Power Management

Currently reads, writes and other ATA commands have direct effects on device power management and the state of the devices rotational media. Since one of the goals of adding NV Cache to a device is to be able to satisfy reads and writes while the rotation media is spun down, there is a need for changes in the power management behavior of a device which supports NV Cache.

Power management behavior would also need to change to express the state of the device's rotational media. This is because one of the most interesting scenarios for NV Cache enabled devices is the ability of the device to complete reads, writes, and other operations while rotational media is spun down.

4 NV Cache Command Proposal

The primary focus of the proposed NV Cache commands is to enable the host to share in the management of the device's NV Cache, which is a new concept to the ATA standard. Consequently many of the concepts that will be used in discussion and in the remainder of this document are also new. This section establishes the goals and requirements for the NV Cache proposal, defines terminology for new NV Cache concepts, and then illustrates the terminology using scenarios that typify NV Cache usage.

4.1 Proposed Solution Requirements

4.1.1 NV Cache Command Requirements

The purpose of NV Cache Commands is to create a mechanism by which the host may manage the contents of the device's NV Cache. It is not required that the host be able to manage every association in the **NV Cache Set**.

Management of the NV Cache must include:

- Method(s) of requesting the HDD to ~~add~~ add or remove LBA(s) from its **NV Cache Pinned Set**.
- Method(s) of determining the LBA(s) currently in the **NV Cache Pinned Set**.
- Method(s) of requesting the HDD use or not use the **NV Cache Unpinned Set** for Power Management purposes.

Managing the NV Cache may not be accomplished with the use of a new or modified version of any ATA read or write command.

The method of transmitting the lists of LBA(s) must have a standard format.

4.1.2 NV Cache Power Management Requirements

The NV Cache Power Management is realized as an input flag to the ATA **Device Configuration Identify Data Structure**. Its purpose is to enable the ATA Power Management state machine to include sensitivity to NV Caches and the device's current **Spindle State**.

The sensitivity exposed by the NV Cache Power Management input is not always desired so it is required that the host has a mechanism to enable and disable this input.

When NV Cache Power Management input is enabled the ATA Power Management must allow the ability to satisfy reads and writes while rotational media is spun down.

There must be a method to query the **Spindle State** to allow the host to make informed decisions about the commands it issues to the device.

4.2 Proposed Definition of Terms for Solution

NV Cache Set The set of LBA(s) currently represented in the device's entire NV Cache.

NV Cache Pinned Set The set of LBA(s) that have been made un-removable from the NV Cache by the host. Writes to LBA(s) represented in the NV Cache Pinned Set always results in valid data in the NV Cache Set.

NV Cache Unpinned Set The set of LBA(s) that are represented in the NV Cache Set but not represented in the NV Cache Pinned Set. The NV Cache Pinned Set and the NV Cache Unpinned Set are mutually exclusive. NV Cache Unpinned Set is completely managed by the device and LBA(s) represented in the NV Cache Unpinned Set can be added or removed from the NV Cache Set at any time.

NV Cache Set Data A data structure representing the standard format of transmitting LBA(s) in the form of a list of LBA Range Entries.

LBA Range Entry A data structure which represents a range of LBA(s). An LBA Range Entry takes the form of an LBA followed by a number which is the number of sequential LBA(s). The LBA is 6 bytes and the range is 2 bytes making every LBA Range Entry 8 bytes in length.

Spindle State The current state of the device's rotational media. There are 2 possible states: spun up/spinning up and spun down/spinning down.

4.3 Solution Overview: Scenarios and Illustrations

4.3.1 NV Cache Sets

The NV Cache is managed as two distinct areas, the NV Cache Pinned Set and the NV Cache Unpinned Set. It is not likely that these areas of the NV Cache will be contiguous. Each LBA stored in the **NV Cache Set** has an attribute which determines if the device may remove the sector from the NV Cache. This 'pinned' attribute represents whether the LBA belongs to the NV Cache Pinned Set or the NV Cache Unpinned Set, and can only be set and cleared by the host.

The host can only manage the **NV Cache Pinned Set**, which is the set of LBAs that the host requires the device to keep in the NV Cache. The remaining NV Cache Set is the **NV Cache Unpinned Set**.

4.3.2 Pinning

Adding or Removing an LBA from the NV Cache Pinned Set is accomplished by setting or clearing the 'pinned' attribute on a mapped sector. If a device's NV Cache Unpinned Set is too full to satisfy an Add request, then the device must remove some or all of the NV Cache Unpinned Set in order to complete the Add request.

When an LBA is pinned, the sector data that is placed into the NV Cache can come from one of two sources: the device's media or the host in the form of a write. The source is determined by the Add command's Populate Immediately (PI) bit.

- If PI is set, the command is not complete until the disk has transferred all of the LBA(s) sector data from the device's media into the NV Cache.
- If PI is not set, the LBA(s) are added to the NV Cache Pinned Set but are marked as containing invalid (stale) data. No sector data is transferred to the NV Cache before the Add command completes.

4.3.2.1 Scenarios

There are two common scenarios where a host will pin an LBA: it wants to pin an LBA for subsequent read operations, or it wants to pin an LBA which is about to be written to.

- If the host knows that the LBA(s) are about to be read, or frequently read from, the host will want to have the sector data come from the magnetic media and will issue the Add command with the PI set.

Example:

- Preloading frequent randomly accessed OS files.
- Preloading boot files in preparation for a system reboot.
- If the host knows that the LBA(s) are about to be written to, the host will want to have the sector data come from the host and will issue the Add command with PI cleared. The device should not implement a background population of the pinned LBA(s) from its media.

Example:

The writing of the hibernate file in preparation for system hibernation.

4.3.3 Cache Management

The NV Cache management refers to the use of ATA commands to query or take action on the contents of the NV Cache Pinned Set. It involves transmitting lists of LBA(s) that need to be used in an NV Cache management action. Some NV Cache management actions take lists of LBA(s) as inputs and some NV Cache management actions give lists of LBA(s) as outputs.

When sending LBA(s) between the host and the device, LBA(s) are grouped into ranges. LBA ranges consist of an initial LBA and a number which indicates the sequential LBA(s) after the initial LBA in the range. More than one LBA range can be sent in a single Add or Remove command. A list of LBA ranges sent in a single Add or Remove command is referred to as the command's NV Cache Set Data.

4.3.3.1 LBA Range Entry

An individual LBA range is called an LBA Range Entry and is represented by 8 bytes. The LBA is expressed by the LBA Range Entry's first 6 bytes and the range length is a zero based number (i.e. 0=0, 1=1, etc.) represented by the remaining 2 bytes. If the 2 byte range length is 0 then the LBA Range Entry is not valid and should be discarded as padding.

Examples:

- If LBA(s) 11, 12, 13, 14, 15, 16, 17, and 18 were in the NV Cache Pinned Set and LBA(s) 10 and 19 were not, LBA(s) 11 through 18 would make one LBA Range Entry which would have the LBA 11 as its first 48 bits and the value of 8 as its next 16 bits. (0000 0000 000B 0008h).
- If only the single LBA 20 was represented in an LBA Range Entry the range value would be 1. (0000 0000 0014 0001h)

The largest range that can be specified in a LBA Range Entry is 65535 and is specified as a range value of all Fs. Multiple LBA Range Entries must be used to specify larger range values.

4.3.3.2 NV Cache Set Data

NV Cache Set Data consists of an ascending ordered, non overlapping series of 8 byte **LBA Range Entries** which have 8 byte alignment within the NV Cache Set Data. An NV Cache Set Data transfer is a multiple of 512 bytes. In the event that not all of the last block's 64 entries are needed, the unused LBA Range Entries must have their LBA and range length set to 0.

During an Add command, the Populate Immediately (PI) bit in the command applies to all LBA(s) specified in the NV Cache Set Data. (See the Pinning section above) [fs5]

4.3.3.3 NV Cache Remaining for Pinned LBA(s)

The NV Cache can run out of space to hold any more pinned LBA(s) data. The number of the remaining available spaces shall be returned to the host during the completion of NV Cache Commands. This is referred to as LBA(s) Remaining. (Please refer to section 7.1..1.4 and section 7.1.2.4 for the detailed return data structure defines.)

LBA(s) Remaining is 6 byte number that always represents the number of LBA(s) in the total NV Cache size minus the number of LBA(s) currently in the Pinned Cache Set.

4.3.3.4 NV Cache Set Data as Input

NV Cache commands that take an NV Cache Set Data as Input can be thought of as the Set operations. Set operations change the members of the NV Cache Pinned Set.

4.3.3.5 NV Cache Set Data as Output

NV Cache commands that return an NV Cache Set Data as Output can be thought of as the GET operations. The GET operations return the members of the current NV Cache Pinned Set or NV Cache misses members.

4.3.4 Behavior over Power Cycle

4.3.4.1 Query NV Cache Misses

In order to gather information about the system's BIOS read and write behavior, a query of NV Cache misses is necessary. Please see section 7.1.4.

4.3.4.2 Rotational Media State after Power Cycle

It is not always desirable to have the device spin up its rotational media following a power ON. Additional power consumption, heat generation and susceptibility to shock may not warrant the spin up at this time if the host will not need to access LBA(s) on the rotational media.

If the device was in NV Cache Power Mode when it was powered down, it should not spin up the rotational media on power up until it receives a request that it can not satisfy without spinning up the rotational media.

4.3.5 Preparing to Pin a Large NV Cache Set Data

Before the pinning of a large number of LBA(s) in the NV Cache Pinned Set it is desirable for the host to instigate a flush of some or all of the NV Cache Unpinned Set. This ensures that the potentially lengthy flushing operation in a large NV Cache is completed as quickly and early as possible. This is accomplished using the Flush NV Cache command described in section 7.1.5.

5 Proposed Changes to ATA8-ACS

The following ATA8-ACS sections require modification:

- The Identify Device command is modified to be able to identify device support for the new command set and power mode.
- The Device Configuration Overlay feature set is used to control device behavior for the NV Cache command set and power mode.
- The Security Feature Set is modified to abort commands in the NV Cache Feature Set when the drive is Locked.
- Long Logical Sector is updated.

5.1 Identify Device

Determining whether a given ATA HDD supports NV Cache, including modes and size of the cache is a straight forward process. Six words are used in the IDENTIFY DEVICE DATA sent by the HDD during drive initialization. For a device where the NV Cache exists, the device must set new bits to tell the host that it can support the NV Cache Modes features.

The following would modify ATA8-ACS IDENTIFY DEVICE data

IDENTIFY DEVICE NV Cache Word

Word	O/M	F/V	Bit	Description
214	O			
		F	12-15	NV Cache Commands Version
		F	8-11	NV Cache Power Modes Version
			5-7	Reserved
		V	4	1 = NV Cache Commands Supported and Enabled
			1-3	Reserved
		V	0	1 = NV Cache Power Modes Supported
215	O			
		V	0-15	NV Cache Size in Logic Block Size <u>L</u> MSW
216	O			
		V	0-15	NV Cache Size in Logic Block Size <u>M</u> LSW
217	O			
		F	0-15	Read Transfer Speed of NV Cache in MB/s
218	O			
		F	0-15	Write Transfer Speed of NV Cache in MB/s
219	O			
			8-15	Reserved
		F	0-7	Device Estimated Time to Spin Up in Seconds
220-221				Reserved

If bit 0 of word 214 is set to 1, the device supports the power modes of the NV Cache feature set. Bits 8 to 11 specify the version of the NV Cache Modes supported.

If bit 4 of word 214 is set to 1, the device supports the commands of the NV Cache feature set. Bits 12 to 15 specify the version of the NV Cache Commands supported.

Both NV Cache power mode version (word 214 bits 8-11) and NV Cache command version (word 214 bits 12-15) support fields shall be set to 0.

Words 215 and 216 specify the maximum number of logical sectors that the device's NV Cache Set can contain for the host to pin.

Word 217 specifies the maximum sustained transfer speed of the device's NV Cache during a Read in megabytes per second.

Word 218 specifies the maximum sustained transfer speed of the device's NV Cache during a Write in megabytes per second.

Word 219 bits 0-7 specifies a value which will be the device's estimate of the amount of time it takes to be able to satisfy a read or write request from its rotational media when the read or write request is received while the rotational media is not spinning.

5.2 Device Configuration Overlay for NV Cache Feature Set

The host enables/disables support for the NV Cache features through the Device Configuration Overlay feature set commands. The Device Configuration Identify Data Structure is modified to add the new NV Cache feature sets. The default setting in such devices and between power cycles is to have this feature set Enabled. The following is added as a description for the new subcommand code.

Device Configuration Identify Data Structure

Word	BIT	Content
...	...	
21	14	1 = Reporting Support for NV Cache Power Management Feature Set is Allowed
21	15	1 = Reporting Support for NV Cache Command Feature Set is Allowed
...	...	

5.3 Security Feature Set

Command	Locked	Unlocked	Frozen
...			
Add LBA(s) to NV Cache Pinned Set	Command aborted	Executable	Executable
Remove LBA(s) From NV Cache Pinned Set	Command aborted	Executable	Executable
Set NV Cache Power Mode	Command aborted	Executable	Executable
Return From NV Cache Power Mode	Command aborted	Executable	Executable
Flush NV Cache	Command aborted	Executable	Executable
Query NV Cache Pinned Set	Command aborted	Executable	Executable
Query NV Cache Misses	Command aborted	Executable	Executable
...			

5.4 Long Logical Sector Feature Set for Non-Packet Devices

Table - Long Logical Sector Function

Command	Words Per Sector Transferred
...	...
Add LBA(s) to NV Cache Pinned Set	256
Remove LBA(s) From NV Cache Pinned Set	256
Set NV Cache Power Mode	-
Return From NV Cache Power Mode	-
Flush NV Cache	-
Query NV Cache Pinned Set	256
Query NV Cache Misses	256
...	...

6 Proposed New NV Cache Power Management for ATA8-ACS

In addition to the definition of new NV Cache Power Management Feature Set commands, the Check Power Mode command is modified to reflect the new NV Cache Power modes.

6.1 NV Cache Power Management Feature Set

The optional NV Cache Power Management feature set permits a host to modify the behavior of a device in a manner that allows the device to improve response times to read and write commands while reducing the device's power consumption.

Commands unique to the NV Cache Power Management feature set use a single command code and are differentiated from one another by the value placed in the Features register. A device that implements the NV Cache Power Management feature set shall implement the following commands:

- Set NV Cache Power Mode
- Return From NV Cache Power Mode

Individual NV Cache Power Management commands are identified by the value placed in the Feature register as shown below.

Value	Command	Data Transfer
00h	Set NV Cache Power Mode Ext	-
01h	Return From NV Cache Power Mode Ext	-
02h-0Fh	Reserved	-
10h-2Fh	Reserved for NV Cache Commands feature set	-
30h-CFh	Reserved	-
D0h-EFh	Vendor specific	-
FFh	Reserved	-

6.1.1 Set NV Cache Power Mode – B6h/00h, Non-data

6.1.1.1 Description

This command is the only way to cause the device to set the NV Cache Power Mode. Upon completion of this command the device shall treat the NV Cache Power Mode value to be TRUE for all Power Management state transitions and ~~should strive to power down the spindle and use the NV Cache to satisfy read and write requests whenever possible~~ in order to reduce device power consumption.

If the device cannot satisfy a read or write request from its NV Cache it shall spin up and stay spun up for at least as many seconds as the value in Sector Count.

6.1.1.2 Inputs

Word	Name	Description
00h	Feature	00h
01h	Count	Bit Description 15:0 Minimum High-Power Time - See clause 6.1.1.3
02h-04h	LBA	Reserved
05h	Command	B6h

6.1.1.3 Minimum High-Power Time

Contains minimum value, in seconds, that the device shall stay in a high power state in the case that the device must enter the high power state to access it's media while NV CACHE Power Mode is set. The high power state can include any ~~PM~~ Power Management state in which the media is spun up and readily available.

This timer can be aborted upon receiving the 'Return From NV Cache Power Mode' command.

The maximum amount of time the device shall keep the media spun up is vendor specific.

6.1.1.4 Normal Outputs

See Table 64.

6.1.1.5 Error Outputs

See Table 78

6.1.2 Return From NV Cache Power Mode – B6h/01h, Non-data

6.1.2.1 Description

This command is the only way to cause the device to clear the NV Cache Power Mode. Upon completion of this command the device shall treat the NV Cache Power Mode value to be FALSE for all Power Management state transitions.[fs6]

6.1.2.2 Inputs

Word	Name	Description
00h	Feature	01h
01h	Count	Reserved
02h-04h	LBA	Reserved
05h	Command	B6h

6.1.2.3 Normal Outputs

See Table 64.

6.1.2.4 Error Outputs

See Table 78

6.2 Proposed Power State Changes

When the NV Cache Power Mode is set, the device implements an aggressive policy to remove power from its rotational media and satisfy all reads and writes from the device's NV Cache. If a device is not capable of satisfying a read or write from its NV Cache it shall service the read or write request through other means.

Since all IO operations happen in PM0:Active mode, it is the only state that requires attention. The only aspect of the NV Cache feature that affects the Power Management feature is the NV Cache Power Mode input. The purpose of the NV Cache Power Mode input is to enable and disable the aggressive spinning down of the device while it is in the PM0:Active mode. However, the Power Management state diagram does not reflect spindle state and the NV Cache Power Mode input does not affect any Power Management state transitions. Consequently no changes to the Power Management state diagram are needed.

The ~~effects of the NV Cache Power Mode input must be surfaced somewhere~~[fs7], ~~even if the Power Management state machine is not the correct location. The affects~~ are as follows:

- When the device is Powered Up, it shall satisfy requests from NV Cache if possible independent of the NV Cache Power Mode. ~~This avoids BIOSes~~[fs8] ~~to avoid from having~~ to be NV Cache sensitive.[fs9]
- When NV Cache Power Mode is cleared and the device is in PM0:Active it shall spin up.[fs10]
- When NV Cache Power Mode is set and the device is in PM0:Active the device shall enact its aggressive policy to remove power from its rotational media. The only requirement of a device's aggressive policy is that when the device is spun up, it remains spun up for at least the amount of time specified in the Set NV Cache Power Mode command. The remaining aspects of the aggressive policy are vendor specific.

- A Device Configuration Overlay command that disables the NV Cache Power Mode support in the device causes the NV Cache Power Mode to be cleared.

6.3 Check Power Mode Changes

Because much of the aggressive policy to remove power from the device's rotational media is vendor specific the Check Power Modes command must be updated to be able to respond with a value that indicates that the device is in the NV Cache Power Mode.

The only change to the Check Power Modes command would be adding that value to the Features values:

Features result value –

- 40h - device is in NV Cache Power Mode and the spindle is spun down or spinning down.
- 41h - device is in NV Cache Power Mode and the spindle is spun up or spinning up.

7 Proposed New NV Cache Commands for ATA8-ACS

The optional NV Cache Commands feature set permits a host to modify the NV Cache Pinned Set of a device in a manner that allows the device to improve response times to read and write commands while reducing the device's power consumption.

The NV Cache Commands feature set provides a set of commands that guide a device's management of the contents of its NV Cache.

Commands unique to the NV Cache Commands feature set use a single command code and are differentiated from one another by the value placed in the Features register. A device that implements the NV Cache Commands feature set shall implement the following commands:

- Add LBA(s) to NV Cache Pinned Set
- Remove LBA(s) From NV Cache Pinned Set
- Query NV Cache Pinned Set
- Query NV Cache Misses
- Flush NV Cache

The NV Cache Commands feature set requires that the device has DMA Ext Command Block registers support

Individual NV Cache Commands are identified by the value placed in the Feature register as shown below.

The Table of Feature Register Values

Value	Command	Data Transfer
00-01h	Reserved for NV Cache Power Management feature	-
02h-0Fh	Reserved	-
10h	Add LBA(s) to NV Cache Pinned Set DMA Ext	Host to Device
11h	Remove LBA(s) From NV Cache Pinned Set DMA Ext	Host to Device
12h	Query NV Cache Pinned Set DMA Ext	Device to Host
13h	Query NV Cache Misses DMA Ext	Device to Host
14h	Flush NV Cache	-
15h-2Fh	Reserved	-
30h-CFh	Reserved	-
D0h-EFh	Vendor specific	-
FFh	Reserved	-

7.1.1 Add LBA(s) to NV Cache Pinned Set – B6h/10h, DMA Ext

7.1.1.1 Description

This command adds the LBA(s) specified in the NV Cache Set Data to the NV Cache Pinned Set if they are not already so.

If the PI bit is set to one the command shall not complete until the NV Cache population is complete. If the PI bit is set to zero, the command shall complete immediately and the population of the sector data shall be completed on subsequent Write operations to that LBA. If a Read operation occurs to this LBA before the sector data is populated in the NV Cache then this data should be sourced from the valid data located on the magnetic media and may require a disk spin up.

If an LBA Range Entry specified in the NV Cache Set Data does exist but is beyond the range of user-accessible LBA(s), the device shall add the LBA(s) to the NV Cache Pinned Set, but continue to fail all reads and writes to the LBA as before.

The response to this command will be the number of sectors that can still be added to the NV Cache's pinned set.

7.1.1.2 Inputs

Word	Name	Description
00h	Feature	10h
01h	Count	Number of 256 word-blocks to be transferred. 0000h specifies that 65,536 blocks are to be transferred
02h-03h	LBA	Reserved
04h	LBA	Bit Description 15:1 Reserved 0 PI (Populate Immediately) - See clause 7.1.1.3
05h	Command	B6h

[fs12]

Word	Name	Description
00h	Feature	10h
01h	Count	Number of 256 word-blocks to be transferred. 0000h specifies that 65,536 blocks are to be transferred
02-04h	LBA	Bit Description 47:1 Reserved 0 PI (Populate Immediately) - See clause 7.1.1.3
05h	Command	B6h

The data block(s) to be transferred is consisted by LBA Range Entry(s). A LBA Range Entry format is defined in Section 4.3.3.1.

7.1.1.3 Populate Immediately

PI (Populate Immediately) shall be set to one to specify that the LBA(s) specified in the Cache Set Data are to be added to the device's NV Pinned Cache Set and populated immediately from the disk before the command completion. PI shall be set to zero to specify that the LBA(s) specified in the NV Cache Set Data are to be added to the device's NV Pinned Cache Set and populated upon subsequent Write operation to that LBA.

7.1.1.4 Normal Outputs

Word	Name	Description
00h	Error	00h
01h	Count	Reserved
02h-04h	LBA	Bit Description 47:0 Unpinned LBA(s) Remaining
05h	Status	Bit Description 7:6 Transport Dependent - See clause 6.2.11 of ATA8-ACS 5 Device Fault - See clause 6.2.4 of ATA8 4 N/A 3 Transport Dependent - See clause 6.2.11 of ATA8-ACS 2 N/A 1 N/A 0 Error - See clause 6.2.3 of ATA8-ACS

7.1.1.5 Error Outputs

Word	Name	Description
00h	Error	Bit Description
		15:8 Reserved
		7:02 Reserved
		1 Insufficient LBA Range Entries remaining.
		0 Insufficient NV Cache space
01h	Count	Reserved
02h-04h	LBA	Reserved
05h	Status	Bit Description
		7:6 Transport Dependent - See clause 6.2.11 of ATA8-ACS
		5 Device Fault - See clause 6.2.4
		4 N/A
		3 Transport Dependent - See clause 6.2.11 of ATA8-ACS
		2 N/A
		1 N/A
0 Error - See clause 6.2.3 of ATA8-ACS		

If the Add command is failed, non of the requested LBA(s) are added to the Pinned Set.

7.1.2 Remove LBA(s) From NV Cache Pinned Set – B6h/11h, DMA Ext

7.1.2.1 Description

This command removes the LBA(s) specified in the NV Cache Set Data from the NV Cache's pinned set, no longer including them in the set of LBA(s) that must always be mapped in the NV Cache.

If the NV Cache Set Data specifies an LBA not in the NV Cache Pinned Set, the LBA shall be ignored without causing an error.

The response to this command will be the number of sectors that the NV Cache can still add to the NV Cache's pinned working set as specified in the Sector Count, LBA Low, LBA Mid and LBA High registers where the LBA High register is the MSB.

7.1.2.2 Inputs

Word	Name	Description
00h	Feature	11h
01h	Count	Number of 256 word-blocks to be transferred. 0000h specifies that 65,536 blocks are to be transferred
02-04h	LBA	Bit Description
		47:1 Reserved 0 UA(Unpin All) - See clause 7.1.2.3
05h	Command	B6h

The data block(s) to be transferred is consisted by LBA Range Entry(s). A LBA Range Entry format is defined in Section 4.3.3.1.

7.1.2.3 Unpin All

UA (Unpin All) indicates that the Value in Sector Count should be ignored, and the NV Cache Pinned Set shall have no LBA(s) mapped.

7.1.2.4 Normal Outputs

Word	Name	Description
00h	Error	00h
01h	Count	Reserved
02h-04h	LBA	Bit Description
		47:0 Unpinned LBA(s) Remaining
05h	Status	Bit Description
		7:6 Transport Dependent - See clause 6.2.11 of ATA8-ACS
		5 Device Fault - See clause 6.2.4 of ATA8
		4 N/A
		3 Transport Dependent - See clause 6.2.11 of ATA8-ACS
		2 N/A
		1 N/A
0 Error - See clause 6.2.3 of ATA8-ACS		

7.1.2.5 Error Outputs

00h	Error	Bit Description
		7:1 Reserved
		0 A Partial Remove of An Existing LBA Range Entry From Pinned Set
01h	Count	Reserved
02h-04h	LBA	Reserved
05h	Status	Bit Description
		7:6 Transport Dependent - See clause 6.2.11 of ATA8-ACS
		5 Device Fault - See clause 6.2.4 of ATA8
		4 N/A
		3 Transport Dependent - See clause 6.2.11 of ATA8-ACS
		2 N/A
		1 N/A
0 Error - See clause 6.2.3 of ATA8-ACS		

If the Remove command is failed, and a subsequence Read command to the same Pinned Set should complete with no error.

7.1.3 Query NV Cache Pinned Set – B6h/12h, DMA Ext

7.1.3.1 Description

This command requests the device to send the LBA Ranges currently in the NV Cache pinned set in 512 bytes block equal to the number in Block Count. The LBA Ranges sent must be in numerical order. If a device does not have as many LBA Ranges as are requested in the transfer, the unused LBA Ranges shall be filled with zero.

7.1.3.2 Inputs

Word	Name	Description
00h	Feature	12h
01h	Count	Number of 256 word-blocks to be transferred. 0000h specifies that 65,536 blocks are to be transferred
02-04h	LBA	Starting LBA
05h	Command	B6h

[fs13]

The data block(s) to be transferred is consisted by LBA Range Entry(s). A LBA Range Entry format is defined in Section 4.3.3.1.

Starting LBA is the first LBA to be queried.

7.1.3.3 Normal Outputs

00h	Error	00h
01h	Count	Reserved
02h-04h	LBA	Bit Description
		47:0 TBD
05h	Status	Bit Description
		7:6 Transport Dependent - See clause 6.2.11 of ATA8-ACS
		5 Device Fault - See clause 6.2.4 of ATA8
		4 N/A
		3 Transport Dependent - See clause 6.2.11 of ATA8-ACS
		2 N/A
		1 N/A
0 Error - See clause 6.2.3 of ATA8-ACS		

7.1.3.4 Error Outputs

00h	Error	Bit Description
		7:1 Reserved
		0 TBD
01h	Count	Reserved
02h-04h	LBA	Reserved
05h	Status	Bit Description
		7:6 Transport Dependent - See clause 6.2.11 of ATA8-ACS
		5 Device Fault - See clause 6.2.4 of ATA8
		4 N/A
		3 Transport Dependent - See clause 6.2.11 of ATA8-ACS
		2 N/A
		1 N/A
0 Error - See clause 6.2.3 of ATA8-ACS		

7.1.4 Query NV Cache Misses – B6h/13h, DMA Ext

7.1.4.1 Description

This command requests the device to report NV Cache Misses in LBA Ranges in a single 512 byte block. The LBA Ranges sent must be in the accessed order. If a device does not have as many LBA Ranges as are requested in the transfer, the unused LBA Ranges shall be filled with zero.

7.1.4.2 Inputs

Word	Name	Description
00h	Feature	13h
01h	Count	01h
02-04h	LBA	Reserved
05h	Command	B6h

7.1.4.3 Normal Outputs

Word	Name	Description
00h	Error	00h
01h	Count	Reserved
02h-04h	LBA	Bit Description
		47:0 TBD
05h	Status	Bit Description
		7:6 Transport Dependent - See clause 6.2.11 of ATA8-ACS
		5 Device Fault - See clause 6.2.4 of ATA8
		4 N/A
		3 Transport Dependent - See clause 6.2.11 of ATA8-ACS
		2 N/A
		1 N/A
0 Error - See clause 6.2.3 of ATA8-ACS		

7.1.4.4 Error Outputs

Word	Name	Description
00h	Error	Bit Description
		7:1 Reserved
		0 TBD
01h	Count	Reserved
02h-04h	LBA	Reserved
05h	Status	Bit Description
		7:6 Transport Dependent - See clause 6.2.11 of ATA8-ACS
		5 Device Fault - See clause 6.2.4 of ATA8
		4 N/A
		3 Transport Dependent - See clause 6.2.11 of ATA8-ACS
		2 N/A
		1 N/A
0 Error - See clause 6.2.3 of ATA8-ACS		

7.1.5 Flush NV Cache – B6h/14h, Non-data

7.1.5.1 Description

The device ensures that it has at least as many LBA(s) as specified in the ‘Minimum Number of LBA(s) to Flush-Make Available’ available to be pinned without requiring a future spin up. If the device must spin up in order to make the required number of LBA(s) available it must do so. The command completes immediately and if necessary the device performs the task of removing LBA(s) from the NV Cache Unpinned Set to provide the capacity requested in a prudent fashion. If the device determines it is unable to complete the operation, it shall report an error.

The output of this command is the number of LBA(s) which have yet to be removed from the NV Cache Unpinned Set to satisfy the ‘Minimum Number of LBA(s) to Make Available’ or completely empty the NV Cache Unpinned Set, whichever is lesser.

7.1.5.2 Inputs

Word	Name	Description
00h	Feature	14h
01h	Count	Reserved
02-04h	LBA	Bit Description
		31:0 Minimum Number of LBA(s) to Flush 47:32 Reserved
05h	Command	B6h

7.1.5.3 Normal Outputs

Word	Name	Description
00h	Error	00h
01h	Count	Reserved
02h-04h	LBA	Bit Description
		47:0 Number of LBA(s) still to be removed
05h	Status	Bit Description
		7:6 Transport Dependent - See clause 6.2.11 of ATA8-ACS
		5 Device Fault - See clause 6.2.4 of ATA8
		4 N/A
		3 Transport Dependent - See clause 6.2.11 of ATA8-ACS
		2 N/A
		1 N/A
0 Error - See clause 6.2.3 of ATA8-ACS		

Number of LBA(s) still to be removed is the number of LBA(s) which have not yet to be removed from the NV Cache Unpinned Set to satisfy the 'Minimum Number of LBA(s) to Flush' or completely empty the NV Cache Unpinned Set, whichever is lesser."

7.1.5.4 Error Outputs

Word	Name	Description
00h	Error	Bit Description
		7:1 Reserved
		0 TBD
01h	Count	Reserved
02h-04h	LBA	Reserved
05h	Status	Bit Description
		7:6 Transport Dependent - See clause 6.2.11 of ATA8-ACS
		5 Device Fault - See clause 6.2.4 of ATA8
		4 N/A
		3 Transport Dependent - See clause 6.2.11 of ATA8-ACS
		2 N/A
		1 Device Fail to Flush Required Minimum Number of LBA(s).
0 Error - See clause 6.2.3 of ATA8-ACS		

