



LEGAL NOTICE:

© **Copyright 2007 - 2018 NVM Express, Inc. ALL RIGHTS RESERVED.**

This NVM Express revision 1.3 technical proposal is proprietary to the NVM Express, Inc. (also referred to as "Company") and/or its successors and assigns.

NOTICE TO USERS WHO ARE NVM EXPRESS, INC. MEMBERS: Members of NVM Express, Inc. have the right to use and implement this NVM Express revision 1.3 technical proposal subject, however, to the Member's continued compliance with the Company's Intellectual Property Policy and Bylaws and the Member's Participation Agreement.

NOTICE TO NON-MEMBERS OF NVM EXPRESS, INC.: If you are not a Member of NVM Express, Inc. and you have obtained a copy of this document, you only have a right to review this document or make reference to or cite this document. Any such references or citations to this document must acknowledge NVM Express, Inc. copyright ownership of this document. The proper copyright citation or reference is as follows: "© **2007 - 2018 NVM Express, Inc. ALL RIGHTS RESERVED.**" When making any such citations or references to this document you are not permitted to revise, alter, modify, make any derivatives of, or otherwise amend the referenced portion of this document in any way without the prior express written permission of NVM Express, Inc. Nothing contained in this document shall be deemed as granting you any kind of license to implement or use this document or the specification described therein, or any of its contents, either expressly or impliedly, or to any intellectual property owned or controlled by NVM Express, Inc., including, without limitation, any trademarks of NVM Express, Inc.

LEGAL DISCLAIMER:

THIS DOCUMENT AND THE INFORMATION CONTAINED HEREIN IS PROVIDED ON AN "**AS IS**" BASIS. TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, NVM EXPRESS, INC. (ALONG WITH THE CONTRIBUTORS TO THIS DOCUMENT) HEREBY DISCLAIM ALL REPRESENTATIONS, WARRANTIES AND/OR COVENANTS, EITHER EXPRESS OR IMPLIED, STATUTORY OR AT COMMON LAW, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE, VALIDITY, AND/OR NONINFRINGEMENT.

All product names, trademarks, registered trademarks, and/or servicemarks may be claimed as the property of their respective owners.

NVM Express Workgroup
c/o VTM Group
3855 SW 153rd Drive
Beaverton, OR 97003 USA
info@nvmexpress.org

NVM Express Technical Proposal for New Feature

Technical Proposal ID	TP 4035 Flush Command Behavior with NSID value of FFFFFFFFh
Change Date	06/11/2018
Builds on Specification	NVM Express 1.3a ,ECN 004, and TP 4005a Namespace Write Protect.

Technical Proposal Author(s)

Name	Company
Michael Allison	Intel
Jonathan Hughes	Intel
Fred Knight	NetApp
Peter Onufryk	Microsemi

This technical proposal defines the behavior of a Flush command when the NSID is set to FFFFFFFFh

Revision History

Revision Date	Change Description
1/17/2018	Initial draft containing the changes specified in the NVMe 1.3 ECN 005 presented at the 1/11/2018 meeting that was requested to be a technical proposal due to adding bit definitions.
1/19/2018	Updated TP#
1/25/2018	Requiring that the support be required in version 1.4+. Removed quotes around the binary values. Used 'set' instead of 'cleared' for a value larger than one bit.
2/7/2018	Add clarification to footnote 3 in reference to the addition of footnote 6 in Figure 182.
2/12/2018	Updated Namespace Identified (NSID) definition to include figure 182 now that a NVM Set command deals with NSID of FFFFFFFFh.
3/8/2018	Adding support for Namespace Write Protect to allow Flush command to a namespace in a write protected state. This is only allowed when TP 4005a Namespace Write Protect is approved that forces all Volatile Write cache data to be flushed as part of entering a write protected state. The support for a Flush command with an NSID of FFFFFFFFh on a write protected namespace enables the host to use a single command to flush all namespaces when namespaces exist in a write protected state.
3/15/2018	Flush command on Name Space Write protect moved to TP 4005a.
6/11/2018	Ratified

<Editor's note: The figure and section numbers used as based on the NVMe 1.3a PDF files available on the NVMe Website. There is issues with the numbering in relation to the WORD document available.>

Description of Specification Changes

Make changes to Figure 11 (Command Format – Admin and NVM Command Set) as shown below:

07:04	<p>Namespace Identifier (NSID): This field specifies the namespace that this command applies to. If the namespace identifier is not used for the command, then this field shall be cleared to 0h. The value FFFFFFFFh in this field is a broadcast value (refer to section 6.1), where the scope (e.g., the NVM subsystem, all attached namespaces, or all namespaces in the NVM subsystem) is dependent on the command. Refer to Figure 41, and Figure 42, and Figure 188 for Admin commands that support the use of the value FFFFFFFFh in this field.</p> <p>Specifying an inactive namespace identifier (refer to section 6.1.4) in a command that uses the namespace identifier shall cause the controller to abort the command with status Invalid Field in Command, unless otherwise specified. Specifying an invalid namespace identifier (refer to section 6.1.2) in a command that uses the namespace identifier shall cause the controller to abort the command with status Invalid Namespace or Format, unless otherwise specified.</p>
-------	---

Make changes to Figure 109 (Identify – Identify Controller Data Structure) as shown below:

525	M	Volatile Write Cache (VWC): This field indicates attributes related to the presence of a volatile write cache in the controller.										
		Bits 7: 4 3 are reserved.										
		Bits 2:1 indicate Flush command behavior (refer to section 6.8) if the NSID value is set to FFFFFFFFh as follows:										
		<table><tr><th>Value</th><th>Definition</th></tr><tr><td>00b</td><td>Support for the NSID field set to FFFFFFFFh is not indicated. Only controllers compliant with versions 1.3 and earlier of the specification shall be allowed to return this value.</td></tr><tr><td>01b</td><td>Reserved.</td></tr><tr><td>10b</td><td>The Flush command does not support the NSID field set to FFFFFFFFh. The controller shall fail a Flush command with the NSID set to FFFFFFFFh with a status code of Invalid Namespace or Format.</td></tr><tr><td>11b</td><td>The Flush command supports the NSID field set to FFFFFFFFh.</td></tr></table>	Value	Definition	00b	Support for the NSID field set to FFFFFFFFh is not indicated. Only controllers compliant with versions 1.3 and earlier of the specification shall be allowed to return this value.	01b	Reserved.	10b	The Flush command does not support the NSID field set to FFFFFFFFh. The controller shall fail a Flush command with the NSID set to FFFFFFFFh with a status code of Invalid Namespace or Format.	11b	The Flush command supports the NSID field set to FFFFFFFFh.
		Value	Definition									
00b	Support for the NSID field set to FFFFFFFFh is not indicated. Only controllers compliant with versions 1.3 and earlier of the specification shall be allowed to return this value.											
01b	Reserved.											
10b	The Flush command does not support the NSID field set to FFFFFFFFh. The controller shall fail a Flush command with the NSID set to FFFFFFFFh with a status code of Invalid Namespace or Format.											
11b	The Flush command supports the NSID field set to FFFFFFFFh.											
Bit 0 if set to '1' indicates that a volatile write cache is present. If cleared to '0', a volatile write cache is not present.												
If a volatile write cache is present, then the host controls whether the volatile write cache is enabled with Set Features specifying the Volatile Write Cache feature identifier (refer to 5.21.1.6). The Flush command (refer to section 6.8) is used to request that the contents of a volatile write cache be made non-volatile.												

Make changes to Figure 188 (Opcodes for NVM Commands) as shown below:

Figure 182: Opcodes for NVM Commands

Opcode by Field			Combined Opcode ²	O/M ¹	Command ³
(07)	(06:02)	(01:00)			
Standard Command	Function	Data Transfer ⁵			
0b	000 00b	00b	00h	M	<u>Flush</u> ⁶
0b	000 00b	01b	01h	M	<u>Write</u>
0b	000 00b	10b	02h	M	<u>Read</u>
0b	000 01b	00b	04h	O	<u>Write Uncorrectable</u>
0b	000 01b	01b	05h	O	<u>Compare</u>
0b	000 10b	00b	08h	O	<u>Write Zeroes</u>
0b	000 10b	01b	09h	O	<u>Dataset Management</u>
0b	000 11b	01b	0Dh	O ⁴	Reservation Register
0b	000 11b	10b	0Eh	O ⁴	Reservation Report
0b	001 00b	01b	11h	O ⁴	Reservation Acquire
0b	001 01b	01b	15h	O ⁴	Reservation Release
Vendor Specific					
1b	na	NOTE 5	80h – FFh	O	Vendor specific
<p>NOTES:</p> <ol style="list-style-type: none"> O/M definition: O = Optional, M = Mandatory. Opcodes not listed are reserved. All NVM commands use the Namespace Identifier field (CDW1.NSID). The value FFFFFFFFh is not supported in this field unless footnote 6 in this figure indicates that a specific command does support that value. Mandatory if reservations are supported as indicated in the Identify Controller data structure. Indicates the data transfer direction of the command. All options to the command shall transfer data as specified or transfer no data. All commands, including vendor specific commands, shall follow this convention: 00b = no data transfer; 01b = host to controller; 10b = controller to host; 11b = bidirectional. This command may support the use of the Namespace Identifier field (CDW1.NSID) set to FFFFFFFFh. 					

Make changes to section 6.8 (Flush command) as shown below:

6.8 Flush command

The Flush command is used to request that the contents of volatile write cache be made non-volatile.

If a volatile write cache is enabled (refer to section 5.21.1.6) then the Flush command shall commit data and metadata associated with the specified namespace(s) to non-volatile media. The flush applies to all commands **for the specified namespace(s) completed by the controller** prior to the submission of the Flush command. The controller may also flush additional data and/or metadata from any namespace.

If bits 2:1 are set to 11b in the VWC field (refer to Figure 109) and the specified NSID is FFFFFFFFh, then the Flush command applies to all namespaces attached to the controller processing the Flush command. If bits 2:1 are set to 10b in the VWC field and the specified NSID is FFFFFFFFh, then the controller fails the command with status code Invalid Namespace or Format. If bits 2:1 are cleared to 00b in the VWC field, then the controller behavior if the specified NSID is FFFFFFFFh is not indicated. Only controllers compliant with versions 1.3 and earlier of the specification shall be allowed to return the value of 00b.

If a volatile write cache is not present or not enabled then, Flush commands shall complete successfully and have no effect.

All command specific fields are reserved.