



LEGAL NOTICE:

© Copyright 2007 - 2021 NVM Express, Inc. ALL RIGHTS RESERVED.

This NVM Express revision 1.4 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 1.4 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 - 2021 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 4097
Change Date	4/29/2021
Builds on Specification	NVM Express 1.4b
References Ratified TPs	

Technical Proposal Author(s)

Name	Company
Fred Knight	NetApp
John Meneghini	NetApp
Murali Rajagopal	VMware
James Smart	Broadcom
David Black	Dell / EMC
Costa Sapuntzakis	Pure Storage

This technical proposal adds clarification for the Admin Abort command and adds a new optional Cancel command.

Revision History

Revision Date	Change Description
9/04/2020	Initial creation.
9/18/2020	Incorporate TWG meeting feedback.
	Switch deferred abort to use ACRE.
10/5/2020	Rework Deferred Abort so no host feature enable is required. Change error code for when CID matches the CID of the Cancel command. Add IOQ deletion/recreate as a possible way to cleanup a hung Cancel command.
10/14/2020	Clarify Barrier requirements (no write modifications after completion of abort/cancel; read may still modify buffers until the command being aborted/cancelled returns status).
10/21/2020	Updates based on 10/20 FMDS meeting and input from David Black and John Meneghini.
12/18/2020	As per TWG discussion, update deferred ABORT description. Remove reference to "candidate" commands.
12/22/2020	Enhance description of Command Abort Requested status.
1/7/2021	Changes from TWG meeting: Change ignored fields to have required values in Cancel Command – 0 for NSID and FFFFh for CID.
1/21/2021	Define "fast abort" and "deferred abort", rewrite barrier description as "no impact" and push any aborted command that did have "impact" into the deferred abort category.
1/24/2021	Use "immediate" and "non-immediate" terms.
1/25/2021	Reorganize some of the text based on review feedback.
1/26/2021	Incorporate comments from FMDS meeting. Incorporate text pushed here from ECN-007.
1/28/2021	Update based on TWG meeting (change non-immediate to deferred; change abort operation to abort action). Reword immediate abort requirements.
2/11/2021	Update text for Immediate Abort requirements in section X.Y. Reword "shall" statements in Cancel command to make it clearer these are "shall TRY" requirements.
2/23/2021	Change Cancel command completion data – change "commands found" to "commands eligible for deferred abort".
3/9/2021	Change NSID in Single Command Cancel action from "shall be 0" to be used and validated (as per footnote 2 in Figure 346). Other editorial clarifications.
3/11/2021	Accept changes based on TWG meeting and send to 30-day review.
4/27/2021	Integrated into the NVMe Base Specification.
4/28/2021	Changed "completion entry" to "completion queue entry". Changed and/or text to use a list to avoid confusion.
4/29/2021	Removed unneeded strikethrough text in second bullet for Command Abort Requested in Figure 128. Removed all comments, accepted all changes, and converted references/cross-references to text.

Black text indicates existing material in the NVMe Base 1.4 specification, Red text indicates material added or modified by this TP. Green text indicates notes.

Incompatible Changes

The Abort command required that for bit 0 of Dword 0 to be cleared to '0' in the CQE of the Abort command, the CQE for the command being aborted be posted prior to the posting of the CQE for the Abort command. To post the CQE for the command being aborted obviously required that all actions associated with that command be stopped. The Abort command is now allowed to post the CQE for the ABORT command and the

CQE for the command being aborted in any order – on the condition that all the actions associated with the command being aborted have stopped.

While the stopping of actions associated with the command being cancelled is compatible, the order of the posting of the CQEs may be considered incompatible.

Description of Specification Changes

Modify the following sections as shown below:

4.6 Completion Queue Entry

...

4.6.1.2.1 Generic Command Status Definition

Figure 128: Status Code – Generic Command Status Values

Value	Description
00h	Successful Completion: The command completed without error.
...	
07h	Command Abort Requested: The command was aborted due to: <ul style="list-style-type: none">an Abort command being received that specified the Submission Queue Identifier and Command Identifier of this command (refer to section 5.1); ora Cancel command being received that specified this command (refer to section 6.TBD).
...	

...

4.6.1.2.2 Command Specific Status Definition

Completion queue entries with a Status Code Type of Command Specific Errors indicate an error that is specific to a particular command opcode. Status codes of 00h to 7Fh are for Admin command errors. Status codes of 80h to BFh are specific to the selected I/O command set.

...

Figure 131: Status Code – Command Specific Status Values, NVM Command Set

Value	Description	Commands Affected
80h	Conflicting Attributes	Dataset Management, Read, Write
81h	Invalid Protection Information	Compare, Read, Verify, Write, Write Zeroes
82h	Attempted Write to Read Only Range	Dataset Management, Write, Write Uncorrectable, Write Zeroes
84h	Invalid Command ID	Cancel
85h to BFh	Reserved	

<Editor's note: The Invalid Command ID error must be valid for all command sets (e.g., NVM, ZNS, KV).>

5.1 Abort command

The Abort command is used to **request an abort of** a specific command previously submitted to the Admin Submission Queue or an I/O Submission Queue. An Abort command ~~is a best effort command; may or may not abort the command that was requested to be aborted.~~ For example, the command to abort may have already completed, currently be in execution, or may be deeply queued.

The controller should process the Abort command as soon as the command is fetched.

To abort a large number of commands (e.g., a larger number of commands than the limit listed in the ACL field), the host should:

- **use the Cancel command (refer to section 6.TBD), if supported, to abort all commands submitted to an I/O Submission Queue or to abort all commands submitted to an I/O Submission Queue for a specific namespace; or**
- **follow the procedures described in section 7.3.3 7.4.3 to delete the I/O Submission Queue and recreate the I/O Submission Queue.**

If an abort action is performed on the command to abort, that action may be:

- **an immediate abort (i.e., the abort action occurs prior to posting the completion queue entry for the Abort command); or**
- **a deferred abort (i.e., the abort action occurs after posting the completion queue entry for the Abort command).**

If an immediate abort is performed, then the controller shall either:

- **post the completion queue entry for the command to abort to the appropriate Admin Completion Queue or I/O Completion Queue with a status code of Command Abort Requested before the completion queue entry for the Abort command is posted to the Admin Completion Queue; or**
- **ensure that there are no subsequent effects of the command to abort, as described in section X.Y, prior to posting the completion queue entry for the Abort command and post the completion queue entry for the command to abort to the appropriate Admin Completion Queue or I/O Completion Queue with a status code of Command Abort Requested. The completion queue entry for the Abort command and the completion queue entry for the command to abort are allowed to be posted in any order.**

If an immediate abort is not performed, then the controller may or may not perform a deferred abort.

The Abort command uses the Command Dword 10 field. All other command specific fields are reserved.

The Abort Command Limit field in the Identify Controller data structure (refer to Figure 249) indicates the controller limit on concurrent execution of Abort commands. A host should not allow the number of outstanding Abort commands to exceed this value. The controller may complete any excess Abort commands with **the status code set to Abort Command Limit Exceeded status.**

5.1.1 Command Completion

Upon completion of the Abort command, the controller posts a completion queue entry to the Admin Completion Queue indicating the status for the Abort command ~~and indicating whether the command to abort was aborted.~~ Dword 0 of the completion queue entry indicates ~~whether information about the command to abort as shown in Figure TBDx. was aborted.~~

Figure TBDx: Abort Command – Completion Queue Entry Dword 0

Bits	Description
31:01	Reserved

Figure TBDx: Abort Command – Completion Queue Entry Dword 0

Bits	Description
00	Immediate Abort Not Performed: Indicates whether or not an immediate abort was performed. If set to '1', then an immediate abort was not performed (i.e., the command to abort was not aborted prior to posting the CQE for the Abort command) for any reason (e.g., the immediate abort requirements were not met, the requested command to abort was not found). If cleared to '0', then an immediate abort was performed (i.e., the command to abort was aborted prior to posting the CQE for the Abort command).

If the Immediate Abort Not Performed bit in Dword 0 (refer to Figure TBDx) is set to '1' in the completion queue entry for the Abort command and a deferred abort is performed, then the controller shall set the status code for the command to abort to Command Abort Requested. The host should examine the status in the completion queue entry of the command to abort to determine if the command was aborted.

~~If the command to abort was successfully aborted, then a completion queue entry for the aborted command shall be posted to the appropriate Admin or I/O Completion Queue with a status of Command Abort Requested before the completion queue entry for the Abort command is posted to the Admin Completion Queue, and bit 0 of Dword 0 shall be cleared to '0' in the completion queue entry for the Abort command. If the command to abort was not aborted for any reason, then bit 0 of Dword 0 shall be set to '1' in the completion queue entry for the Abort command. If the command to abort was not successfully aborted, for any reason, then bit 0 of Dword 0 shall be set to '1' in the completion queue entry for the Abort command.~~

Command specific status ~~code~~ values associated with the Abort command are defined in Figure 144.

...

<Editor's Note: New section - Suggest 5.1.2 (right after ABORT completion section.)>

X.Y Immediate Abort requirements

If the controller performs an immediate abort, then the controller shall ensure that, other than the posting of the completion queue entry for the command to abort, there are no effects of that command (e.g., no host memory access, command processing does not modify: NVM media, NVM Set state, Endurance Group state, namespace state, controller state, domain state, or NVM subsystem state) subsequent to the posting of the CQE for the command that requested the abort. If the controller is not able to ensure there are no effects of the command to abort subsequent to the posting of the CQE for the command that requested the abort, then the controller shall not perform an immediate abort on that command.

For example, if a command:

- requests a data transfer;
- the data transfer has been initiated and not completed; and
- the CQE for that command was not posted prior to the posting of the CQE for the command that requested the abort,

then the controller is prohibited from performing an immediate abort.

Performing an immediate abort effects the information returned in the CQE for the command that requested the abort as described in section 5.1.1 for an Abort command and in section **6.TBD.1** for a Cancel command.

5.15 Identify command

...

5.15.1.1 Identify Controller data structure (CNS 01h)

...

Figure 251: Identify – Identify Controller Data Structure

Bytes	O/M ¹	Description
...		
258	M	Abort Command Limit (ACL): This field is used to convey indicates the maximum number of concurrently executing Abort commands (refer to section 5.1) on the Admin Queue supported by the controller (refer to section 5.1). This is a 0's based value. It is recommended that implementations support concurrent execution of a minimum of four Abort commands.
...		

6 NVM Command Set

...

Figure 350: Opcodes for NVM Commands

Opcode by Field			Combined Opcode ¹	Command ²	Reference Section
(07)	(06:02)	(01:00)			
Standard Command	Function	Data Transfer ³			
0b	000 00b	00b	00h	Flush ⁴	6.8
0b	000 00b	01b	01h	Write	6.15
0b	000 00b	10b	02h	Read	6.9
0b	000 01b	00b	04h	Write Uncorrectable	6.16
0b	000 01b	01b	05h	Compare	6.6
0b	000 10b	00b	08h	Write Zeroes	6.17
0b	000 10b	01b	09h	Dataset Management	6.7
0b	000 11b	00b	0Ch	Verify	6.14
0b	000 11b	01b	0Dh	Reservation Register	6.11
0b	000 11b	10b	0Eh	Reservation Report	6.13
0b	001 00b	01b	11h	Reservation Acquire	6.10
0b	001 01b	01b	15h	Reservation Release	6.12
0b	001 10b	00b	18h	Cancel ⁴	6.TBD
Vendor Specific					
1b	n/a	NOTE 3	80h to FFh	Vendor specific	
NOTES: 1. Opcodes not listed are reserved. 2. All NVM commands use the Namespace Identifier (NSID) field. The value FFFFFFFFh is not supported in this field unless footnote 4 in this figure indicates that a specific command does support that value. 3. 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. 4. This command may support the use of the Namespace Identifier (NSID) field set to FFFFFFFFh.					

<Editor's note: The op-code assigned must not overlap any existing command set op-codes. This command is usable on all command sets.>

Add the following sections alphabetically (before the Compare command) as shown below:

6.TBD Cancel command

The Cancel command is used to request an abort for specified commands submitted on the same I/O Submission Queue to which the Cancel command is submitted. The Cancel command may apply to a single namespace or to multiple namespaces. The Cancel command may be deeply queued. Some of the commands to abort may have already completed, currently be processing, or be deeply queued. As a result, the performance of the Cancel command may be impacted.

To abort an Admin command, the host uses the Abort command (refer to section 5.1). To abort a Cancel command, the host may:

- use the Abort command;
- use a second Cancel command with a Cancel Action set to Single Command Cancel (i.e., 00b) that specifies the CID of the Cancel command to abort; or
- follow the procedures described in section 7.4.3 to delete the I/O Submission Queue and recreate the I/O Submission Queue.

The controller applies the specified Cancel Action (refer to Figure TBD2) to all outstanding commands that have been fetched prior to the processing of the Cancel command. The controller may or may not apply the specified Cancel Action to commands that are fetched:

- after the start of processing of the Cancel command; and
- before the completion of the Cancel command.

There is no requirement to apply the specified Cancel Action to commands that have not been fetched by the controller.

If an abort action is performed on a command to abort, that action may be:

- an immediate abort (i.e., the abort action occurs prior to posting the completion queue entry for the Cancel command); or
- a deferred abort (i.e., the abort action may or may not occur prior to posting the completion queue entry for the Cancel command).

For each command on which an immediate abort is performed, the controller shall meet the immediate abort requirements (refer to section X.Y) for that command before posting the completion queue entry for the Cancel command.

For each command on which a deferred abort is performed, there are no requirements on the ordering of posting of the completion queue entry for the Cancel command (refer to section 6.TBD.1).

For each command on which an abort action is performed (i.e., immediate abort or deferred abort) the status code shall be set to Command Abort Requested in the completion queue entry for that command.

If the Cancel command is supported, then the Commands Supported and Effects log page shall be supported.

The Cancel command uses the Command Dword 10 and Command Dword 11 fields. All other command specific fields are reserved.

Figure TBD1: Cancel – Command Dword 10

Bits	Description
31:16	<p>Command Identifier (CID): This field specifies the command identifier of the command to be aborted (i.e., the CDW0.CID field within the command to be aborted).</p> <p>If the Action Code field is set to Single Command Cancel and this field is set to the CID for this Cancel command, then the controller shall abort the command with a status code of Invalid Command ID.</p> <p>If the Action Code field is set to Multiple Command Cancel and this field is not set to FFFFh, then the controller shall abort the command with a status code of Invalid Field in Command.</p>
15:00	<p>Submission Queue Identifier (SQID): This field specifies the identifier of the Submission Queue that the Cancel command is associated with.</p> <p>If the SQID does not match the Submission Queue Identifier of the submission queue to which the Cancel command is submitted, then the controller shall abort the command with a status code of Invalid Field in Command.</p>

Figure TBD2: Cancel – Command Dword 11

Bits	Description								
31:02	Reserved								
01:00	Action Code:								
	<table><tr><th>Value</th><th>Cancel Action</th></tr><tr><td>00b</td><td><p>Single Command Cancel: The hosts requests that the controller abort the specific command submitted to the specified namespace with the specified CID.</p><p>If the NSID field is set to FFFFFFFFh then, the host requests that the controller abort the specific command submitted to any NSID with the specified CID.</p><p>If the NSID is not set to FFFFFFFFh and the specified CID is not associated with the specified NSID then, the controller shall abort the Cancel command with a status code of Invalid Field in Command.</p><p>If the specified CID is not found, then the controller shall complete the Cancel command with the Commands Eligible for Deferred Abort field cleared to 0h and the Commands Aborted field cleared to 0h.</p></td></tr><tr><td>01b</td><td><p>Multiple Command Cancel: The hosts requests that the controller abort all commands, other than this Cancel command, on the I/O Submission Queue to which the Cancel command was submitted that were submitted to the specified NSID.</p><p>If the NSID is set to FFFFFFFFh then the host requests that the controller abort all commands, other than this Cancel command, submitted to the I/O Submission Queue to which the Cancel command was submitted for all namespaces.</p></td></tr><tr><td>All others</td><td>Reserved</td></tr></table>	Value	Cancel Action	00b	<p>Single Command Cancel: The hosts requests that the controller abort the specific command submitted to the specified namespace with the specified CID.</p> <p>If the NSID field is set to FFFFFFFFh then, the host requests that the controller abort the specific command submitted to any NSID with the specified CID.</p> <p>If the NSID is not set to FFFFFFFFh and the specified CID is not associated with the specified NSID then, the controller shall abort the Cancel command with a status code of Invalid Field in Command.</p> <p>If the specified CID is not found, then the controller shall complete the Cancel command with the Commands Eligible for Deferred Abort field cleared to 0h and the Commands Aborted field cleared to 0h.</p>	01b	<p>Multiple Command Cancel: The hosts requests that the controller abort all commands, other than this Cancel command, on the I/O Submission Queue to which the Cancel command was submitted that were submitted to the specified NSID.</p> <p>If the NSID is set to FFFFFFFFh then the host requests that the controller abort all commands, other than this Cancel command, submitted to the I/O Submission Queue to which the Cancel command was submitted for all namespaces.</p>	All others	Reserved
	Value	Cancel Action							
	00b	<p>Single Command Cancel: The hosts requests that the controller abort the specific command submitted to the specified namespace with the specified CID.</p> <p>If the NSID field is set to FFFFFFFFh then, the host requests that the controller abort the specific command submitted to any NSID with the specified CID.</p> <p>If the NSID is not set to FFFFFFFFh and the specified CID is not associated with the specified NSID then, the controller shall abort the Cancel command with a status code of Invalid Field in Command.</p> <p>If the specified CID is not found, then the controller shall complete the Cancel command with the Commands Eligible for Deferred Abort field cleared to 0h and the Commands Aborted field cleared to 0h.</p>							
	01b	<p>Multiple Command Cancel: The hosts requests that the controller abort all commands, other than this Cancel command, on the I/O Submission Queue to which the Cancel command was submitted that were submitted to the specified NSID.</p> <p>If the NSID is set to FFFFFFFFh then the host requests that the controller abort all commands, other than this Cancel command, submitted to the I/O Submission Queue to which the Cancel command was submitted for all namespaces.</p>							
All others	Reserved								

6.TBD.1 Command Completion

Upon completion of the Cancel command, the controller posts a completion queue entry to the I/O Completion Queue indicating the status for the Cancel command.

If the Cancel Action (refer to Figure TBD2) specified a Single Command Cancel Action and the Commands Aborted field is cleared to 0h, then the host should examine the status in the completion queue entry of the command to abort to determine whether the command was aborted or not (i.e., whether a deferred abort was performed or not).

If the Cancel Action specified a Multiple Command Cancel Action, then the host should examine the status in the completion queue entry of each command to abort to determine whether the command was aborted or not.

Cancel command specific status code values are defined in Figure TBD3.

Figure TBD3: Cancel – Command Specific Status Values

Value	Description
84h	Invalid Command ID: The specified CID matched the CID of this Cancel command.

Dword 0 of the completion queue entry contains information about the number of commands that were aborted by this command. Dword 0 of the completion queue entry is defined in Figure TBD4.

Figure TBD4: Cancel – Completion Queue Entry Dword 0

Bits	Description
31:16	<p>Commands Eligible for Deferred Abort: This field indicates the number of commands that match the specified criteria (refer to Figure TBD1 and Figure TBD2) on which a deferred abort may be performed. A value of 0h indicates that no commands are eligible for a deferred abort or that the controller does not support deferred aborts. A value of FFFFh indicates that FFFFh or more commands are eligible for a deferred abort.</p> <p>If the Cancel Action (refer to Figure TBD2) specified a Single Command Cancel and an immediate abort was performed on the specified command, then this field shall be cleared to 0h.</p>
15:00	<p>Commands Aborted: This field indicates the number of commands on which the controller performed an immediate abort as a result of processing this Cancel command. A value of 0h indicates that the controller did not perform an immediate abort on any commands as a result of processing this Cancel command.</p> <p>If the Cancel Action (refer to Figure TBD2) specified a Single Command Cancel and an immediate abort was performed on the specified command, then this field shall be set to 1h.</p>

...

7 Controller Architecture

...

7.1.1 I/O Controller

...

Figure 423: I/O Controller – NVM Command Set Support

Command	Command Support Requirements ¹
...	
Write Uncorrectable	O
Cancel	O
Compare	O
...	O
Notes:	
1. O = Optional, M = Mandatory, P = Prohibited	
2. Mandatory if reservations are supported as indicated in the Identify Controller data structure.	

...

7.4.3 Queue Abort

To abort a large number of commands, ~~the recommended procedure is to~~ the host may use:

- the Cancel command (refer to section 6.TBD); or
- delete and recreate the I/O Submission Queue (refer to section 7.3.3).

Specifically, to abort all commands that are submitted to ~~the~~ an I/O Submission Queue host software should:

- issue a Cancel command to that queue with the Cancel Action set to Multiple Command Cancel and the NSID field set to FFFFFFFFh; or
- issue a Delete I/O Submission Queue command for that queue. After ~~the~~ that submission queue has been successfully deleted, indicating that all commands have been completed or aborted, then host software should recreate the queue by submitting a Create I/O Submission Queue command. Host software may then re-submit commands to the associated I/O Submission Queue.