



**LEGAL NOTICE:**

© Copyright 2007 to 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 to 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.

The NVM Express® design mark is a registered trademark of NVM Express, Inc.

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

*Technical input submitted to the NVM Express™ Workgroup is subject to the terms of the NVM Express™ Participant's agreement. Copyright © 2014 to 2021 NVMe™ Corporation.*

### NVM Express Technical Proposal for New Feature

Technical Proposal ID	4047b
Change Date	2021-02-08
Builds on Specification	NVM Express Base Specification version 1.4 NVM Express Over Fabrics version 1.1 NVM Express Management Interface Specification Revision 1.1b

#### Technical Proposal Author(s)

Name	Company
Yoni Shternhell	WDC
Judy Brock	Samsung

This technical proposal adds a list of allowed commands that may be processed during a Format NVM operation and defines the appropriate status code to be returned to notify that Admin command are aborted.

#### Revision History

Revision Date	Change Description
2019-05-28	Initial version
2019-06-12	Added new allowed Admin Commands Added new text in the Format In Progress status code
2019-06-24	Remove text from Format In Progress status code. Move the clarification on status codes into a new TP for Multiple Command Sets.
2019-08-22	Added group comments: <ul style="list-style-type: none"><li>• Adding DST and Persistent Event Log as Allowed commands</li></ul>
2019-08-28	Merging Sanitize Operations and Format NVM Command into one Admin Commands Allowed table
2019-09-04	Creating new section under the Admin Command section for allowed commands listing. Both Sanitize Operation and Format NVM command would reference the new section.
2019-09-12	Moved the new text to the introductory of section 5. Text modifications according to the group input
2019-09-16	Updated the introductory text
2019-09-26	Editorial changes part of Phase 3 comments
2019-10-11	Add Telemetry Log to the allowed admin command
2019-11-20	Rewording of the wording on Format NVM command is in progress and Admin commands being processed
2019-12-09	Opcode 7Fh is called out in the NVMe-oF section until the reactor is completed
2020-01-06	Added the NVMe-oF changes.
2020-01-07	Added the Disconnect Command to Fabric Commands allowed table
2020-04-07	Integrated into NVM Express Base Specification version 1.4 and NVM Express Over Fabrics version 1.1

Technical input submitted to the NVM Express™ Workgroup is subject to the terms of the NVM Express™ Participant's agreement. Copyright © 2014 to 2021 NVMe™ Corporation.

2020-04-09	Removed comments. Accepted changes. Ready for ratification.
4047a	
2020-04-17	Apply DST restriction to both Format command and Sanitize operation
2020-04-30	Change "Controller DST" with DST with a NSID set to 0h"
2020-08-07	<ul style="list-style-type: none"> <li>• Device Self-test - additional "Prohibited" restrictions for Sanitize Operation</li> <li>• Get Log Page: <ul style="list-style-type: none"> <li>◦ Sanitize Status was added</li> <li>◦ Vendor Specific and Persistent Event Log are both prohibited for Sanitize.</li> </ul> </li> <li>• Replacing the "N/A" to Prohibited for Sanitize command.</li> </ul>
2020-08-27	<ul style="list-style-type: none"> <li>• Change 'zero' to '0h' in the restriction for the Error Information log page</li> </ul>
2020-10-18	<ul style="list-style-type: none"> <li>• Integrated into the NVMe Express Base Specification.</li> </ul>
2020-10-21	<ul style="list-style-type: none"> <li>• Fixed grammar from technical writer suggestion.</li> </ul>
2020-10-22	<ul style="list-style-type: none"> <li>• Accepted all changes and removed all comments.</li> </ul>
4047b	
2020-11-08	<ul style="list-style-type: none"> <li>• Added changes to NVMe Express Management Interface Specification to indicate which command messages are allowed through OOB interface while Format NVMe command is in progress.</li> </ul>
2020-11-23	<ul style="list-style-type: none"> <li>• Changed "for a Format NVMe command" to "while a Format NVMe command is in progress"</li> <li>• combined the Sanitize and Format NVMe columns in Figure 116 since the contents are the same.</li> </ul>
2020-11-30	<ul style="list-style-type: none"> <li>• integrate comments fr. 11/23 WG meeting</li> </ul>
2020-12-07	<ul style="list-style-type: none"> <li>• accept all changes, change filename to Phase 3</li> </ul>
2021-01-25	<ul style="list-style-type: none"> <li>• ready for integration</li> </ul>
2021-01-31	<ul style="list-style-type: none"> <li>• Integrated into the NVMe Base Specification, the NVMe-Management Interface Specification, Revision 1.1, and the NVMe-Over Fabrics Specification, version 1.1.</li> </ul>
2021-02-01	<ul style="list-style-type: none"> <li>• Change case in the Discussion. Clarified the sentence in section 5.23 of the NVMe Base Specification</li> </ul>
2021-02-04	<ul style="list-style-type: none"> <li>• Accepted changes and removed comments. Ready for integration again.</li> </ul>
2021-02-04	<ul style="list-style-type: none"> <li>• Integrated into the NVMe Base Specification, the NVMe-Management Interface Specification, Revision 1.1, and the NVMe-Over Fabrics Specification, version 1.1, where applicable.</li> </ul>
2021-02-08	<ul style="list-style-type: none"> <li>• Accepted all changes for ratification.</li> </ul>

## Discussion

This technical proposal addresses following issues:

1. The existing NVMe 1.4 specification specifies that I/O commands for a namespace that has a Format NVMe command in progress may abort with a status code of Format in Progress. In some implementations, different commands (e.g. security, Firmware Download, etc.) could be issued while ongoing format operation. For example, security commands in these implementations, are not considered as pure I/O commands or Admin commands and can initiate several internal I/O flows that conflicts with a Format NVMe operation: security erase, access to shadow space/RPMB, etc. Currently

*Technical input submitted to the NVMe Express™ Workgroup is subject to the terms of the NVMe Express™ Participant's agreement. Copyright © 2014 to 2021 NVMe™ Corporation.*

there is no standard way for the controller to abort non-I/O commands during Format NVM command in progress.

**Description for NVMe Changes Document**

1. Format NVM command extended to include a list of allowed commands that may be processed during a Format NVM operation

## Description of the NVM Express Base Specification Revision 1.4 Changes

**Modify portions of section 5.23 (5.23 Format NVM command – NVM Command Set Specific) as shown below:**

...

The Format NVM command shall fail if the controller is in an invalid security state (refer to the appropriate security specification, e.g., TCG Storage Interface Interactions specification). The Format NVM command may fail if there are outstanding I/O commands to the namespace specified to be formatted. I/O commands for a namespace that has a Format NVM command in progress may be aborted and if aborted, the controller should return a status code of Format in Progress. **Refer to section 5 for further information about restrictions on Admin Commands during the processing of a Format NVM command.**

**Modify portions of section 8.15 (8.15 Sanitize Operations) as shown below:**

...

To start a sanitize operation, the host submits a Sanitize command specifying one of the sanitize operation types (i.e., Block Erase, Overwrite, or Crypto Erase). The host sets command parameters, including the Allow Unrestricted Sanitize Exit bit and the No Deallocate After Sanitize bit. After validating the Sanitize command parameters, the controller starts the sanitize operation in the background, updates the Sanitize Status log page and then completes the Sanitize command with Successful Completion status. If the sanitize operation is to be followed by an associated additional media modification operation (refer to NODMMAS in Figure 250), then the associated additional media modification operation shall be completed before the controller reports sanitize operation complete. If a Sanitize command is completed with any status other than Successful Completion, then the controller shall not start the sanitize operation and shall not update the Sanitize Status log page. The controller ignores Critical Warning(s) in the SMART / Health Information log page (e.g., read only mode) and attempts to complete the sanitize operation requested. **Refer to section 5 for further information about restrictions on Admin Commands during a sanitize operation. While a sanitize operation is in progress, all controllers shall abort any commands not listed in Figure 485 with a status of Sanitize In Progress (refer to section 8.15.1).**

...

If a sanitize operation fails, all controllers in the NVM subsystem shall abort any command not allowed during a sanitize operation with a status of Sanitize Failed (refer to section ~~refer to section 8.15.18.15.1~~) until a subsequent sanitize operation is started or successful recovery from the failed sanitize operation occurs.

**Modify portions of section 8.15.1 (8.15.1 Sanitize Command Restrictions) as shown below:**

...

While a sanitize operation is in progress:

- All controllers in the NVM subsystem shall only process the Admin commands listed in **Figure 485** **Figure TBD** subject to the additional restrictions stated in that figure;
- All I/O Commands shall be aborted with a status of Sanitize In Progress;
- Any command or command option that is not explicitly permitted in **Figure TBD** **Figure 485** shall be aborted with a status of Sanitize In Progress if fetched by any controller in the NVM subsystem; and
- The Persistent Memory Region shall be prevented from being enabled (i.e., setting PMRCTL.EN to '1' does not result in PMRSTS.NRDY being cleared to '0').

While a failed sanitize operation has occurred, a subsequent sanitize operation has not started and successful recovery from the failed sanitize operation has not occurred:

*Technical input submitted to the NVM Express™ Workgroup is subject to the terms of the NVM Express™ Participant's agreement. Copyright © 2014 to 2021 NVMe™ Corporation.*

- All controllers in the NVM subsystem shall only process the Sanitize command (refer to section 5.24) and the Admin commands listed in **Figure TBD Figure 485** subject to the additional restrictions noted in that figure;
- All I/O Commands shall be aborted with a status of Sanitize Failed;
- The Sanitize command is permitted with action restrictions (refer to section 5.24);
- Aside from the Sanitize command, any other command or command option that is not explicitly permitted in **Figure TBD Figure 485** shall be aborted with a status of Sanitize Failed if fetched by any controller in the NVM subsystem; and
- The Persistent Memory Region shall be prevented from being enabled (i.e., setting PMRCTL.EN to '1' does not result in PMRSTS.NRDY being cleared to '0').

**Figure 485: Sanitize Operations—Admin Commands Allowed**

Admin Command	Additional Restrictions														
Abort															
Asynchronous Event Request															
Create I/O Completion Queue															
Create I/O Submission Queue															
Delete I/O Completion Queue															
Delete I/O Submission Queue															
Get Features															
Get-Log-Page	The log-pages allowed are listed below:														
	<table><tr><th>Log Pages</th><th>Additional Restrictions</th></tr><tr><td>Error Information</td><td>Return zeroes in the LBA field.</td></tr><tr><td>SMART / Health Information</td><td></td></tr><tr><td>Changed Namespace List</td><td></td></tr><tr><td>Reservation Notification</td><td></td></tr><tr><td>Sanitize Status</td><td></td></tr><tr><td>Asymmetric Namespace Access</td><td></td></tr></table>	Log Pages	Additional Restrictions	Error Information	Return zeroes in the LBA field.	SMART / Health Information		Changed Namespace List		Reservation Notification		Sanitize Status		Asymmetric Namespace Access	
	Log Pages	Additional Restrictions													
	Error Information	Return zeroes in the LBA field.													
	SMART / Health Information														
	Changed Namespace List														
	Reservation Notification														
Sanitize Status															
Asymmetric Namespace Access															
Identify															
Keep-Alive															
NVMe-MI Receive	Prohibited unless explicitly allowed in the NVMe Express Management Interface Specification.														
NVMe-MI Send															
Set Features	Namespace Write Protection Config Feature is not allowed.														
Opcode 7Fh	The Fabric Commands allowed are listed below. Refer to the NVMe over Fabrics specification.														
	<table><tr><th>Fabrics Commands</th><th>Additional Restrictions</th></tr><tr><td>Property Set</td><td></td></tr><tr><td>Connect</td><td></td></tr><tr><td>Property Get</td><td></td></tr><tr><td>Authentication Send</td><td></td></tr><tr><td>Authentication Receive</td><td></td></tr><tr><td>Vendor Specific</td><td>Commands are allowed that do not affect or retrieve user data.</td></tr></table>	Fabrics Commands	Additional Restrictions	Property Set		Connect		Property Get		Authentication Send		Authentication Receive		Vendor Specific	Commands are allowed that do not affect or retrieve user data.
	Fabrics Commands	Additional Restrictions													
	Property Set														
	Connect														
	Property Get														
	Authentication Send														
Authentication Receive															
Vendor Specific	Commands are allowed that do not affect or retrieve user data.														
Vendor Specific	Commands are allowed that do not affect or retrieve user data.														

**Modify portions of section 5 (Admin Command Set) as shown below**

...

Technical input submitted to the NVMe Express™ Workgroup is subject to the terms of the NVMe Express™ Participant's agreement. Copyright © 2014 to 2021 NVMe™ Corporation.

**Figure 143: Opcodes for Admin Commands – NVM Command Set Specific**

Opcode (07)	Opcode (06:02)	Opcode (01:00)	Opcode <sup>1</sup>	Namespace Identifier Used <sup>2</sup>	Command
Generic Command	Function	Data Transfer <sup>3</sup>			
1b	000 00b	00b	80h	Yes	Format NVM
1b	000 00b	01b	81h	NOTE 4	Security Send
1b	000 00b	10b	82h	NOTE 4	Security Receive
1b	000 01b	00b	84h	No	Sanitize
1b	000 01b	10b	86h	NOTE 5	Get LBA Status

NOTES:

1. NVM Command Set Specific opcodes not listed are reserved.
2. A subset of commands use the Namespace Identifier (NSID) field. If the Namespace Identifier field is used, then unless otherwise specified, the value FFFFFFFFh is supported in this field. When this field is not used, the field is cleared to 0h as described in Figure 108.
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. The use of the Namespace Identifier is Security Protocol specific.
5. This command does not support the use of the Namespace Identifier (NSID) field set to FFFFFFFFh.

Formatted: Default Paragraph Font

Formatted: Default Paragraph Font

Formatted: Default Paragraph Font

Formatted: Default Paragraph Font

Formatted: Default Paragraph Font

Figure TBD lists the Admin commands that are allowed during the processing of a sanitize operation and a recommended list of the Admin commands to that should be allowed for during the processing of a Format NVM command.

If a Format NVM command is in progress, then an Admin command not listed in Figure TBD that is submitted for any namespace affected by that Format NVM command may be aborted. If aborted for that reason, then a status code of Format in Progress should be returned.

If there are Admin commands not listed in Figure TBD being processed for a namespace, then a Format NVM command which is submitted that affects that namespace may be aborted. If aborted for that reason, then a status code of Command Sequence Error should be returned.

**Figure TBD: Sanitize Operations and Format NVM Command – Admin Commands Allowed**

Admin Command	Additional Restrictions for Format NVM command	Additional Restrictions for sanitize operations
Abort		
Asynchronous Event Request		
Create I/O Completion Queue		
Create I/O Submission Queue		
Device Self-test	Only Controller DST should be allowed	Prohibited
Delete I/O Completion Queue		
Delete I/O Submission Queue		
Get Features		

Technical input submitted to the NVM Express™ Workgroup is subject to the terms of the NVM Express™ Participant's agreement. Copyright © 2014 to 2021 NVMe™ Corporation.

**Figure TBD: Sanitize Operations and Format NVM Command – Admin Commands Allowed**

Admin Command	Additional Restrictions for Format NVM command	Additional Restrictions for sanitize operations
Get Log Page	The log pages allowed for both Format NVM command and sanitize operations are listed below.	
	<b>Log Pages</b>	<b>Additional Restrictions for both Format NVM command and sanitize operations</b>
	Error Information	Return 0h in the LBA field.
	SMART / Health Information	
	Changed Namespace List	
	Reservation Notification	
	Asymmetric Namespace Access	
	Sanitize Status	
	Vendor specific	Prohibited for Sanitize
	Persistent Event Log	Prohibited for Sanitize
Identify		
Keep Alive		
NVMe-MI Receive	Allowed unless explicitly prohibited in the NVMe Express Management Interface Specification.	
NVMe-MI Send		
Sanitize		Prohibited
Set Features	Namespace Write Protection Config Feature is not allowed.	
Opcode 7Fh	The Fabric Commands allowed are listed below. Refer to the NVMe over Fabrics specification.	
	<b>Fabrics Commands</b>	<b>Additional Restrictions for both Format NVM command and sanitize operations</b>
	Property Set	
	Connect	
	Disconnect	
	Property Get	
	Authentication Send	
	Authentication Receive	
	Vendor Specific	Commands are allowed that do not affect or retrieve user data.
Vendor Specific	Commands are allowed that do not affect or retrieve user data.	



## Description of the NVMe Express Over Fabrics Specification Revision 1.1 Changes

Modify portions of section 3 as shown below:

### 3 Commands

Fabrics commands are used to create queues and initialize a controller. Fabrics commands have an Opcode field of 7Fh. Fabrics commands are processed regardless of the state of controller enable (CC.EN). The Fabrics command capsule is defined in section 2.1 and the Fabrics response capsule and status is defined in section 2.2.

Restrictions on processing commands listed in Figure 14 are defined in the Admin Command Set in section 5 of the NVMe Base specification (e.g., while the NVMe subsystem is performing a sanitize operation or processing a Format NVM command).

Figure 14: Fabrics Command Types

Command Type by Field			Combined Command Type <sup>2</sup>	O/M <sup>1</sup>	I/O Queue <sup>3</sup>	Command
(07)	(06:02)	(01:00)				
Generic Command	Function	Data Transfer <sup>4</sup>				
0b	000 00b	00b	00h	M	No	Property Set
0b	000 00b	01b	01h	M	Yes	Connect <sup>5</sup>
0b	000 01b	00b	04h	M	No	Property Get
0b	000 01b	01b	05h	O	Yes	Authentication Send
0b	000 01b	10b	06h	O	Yes	Authentication Receive
0b	000 10b	00b	08h	O	Yes	Disconnect
Vendor Specific						
1b	na	na	C0h to FFh	O		Vendor specific
NOTES:						
1. O/M definition: O = Optional, M = Mandatory.						
2. Opcodes not listed are reserved.						
3. All Fabrics commands, other than the Disconnect command, may be submitted on the Admin Queue. The I/O Queue supports Fabrics commands as specified in this column. If a Fabrics command that is not supported on an I/O Queue is sent on an I/O Queue, that command shall be aborted with a status code of Invalid Field in Command.						
4. 00b = no data transfer; 01b = host to controller; 10b = controller to host; 11b = reserved						
5. The Connect command is submitted and completed on the same queue that the Connect command creates. Refer to section 1.5.7.						

## Description of the NVMe Express Management Interface Specification Revision 1.1b Changes

Modify Section 6.3 (Sanitize Operation) as follows:

### 6.3 Sanitize Operation and Format NVM Command

Figure 116 specifies the Command Messages allowed during a sanitize operation and the Command Messages that should be allowed during the processing of a Format NVM command. Refer to the NVMe Express specification for the definition of a sanitize operation.

**Figure 116: Command Messages Allowed During Sanitize Operation and During the Processing of a Format NVM Command**

Command Set	Command Message	Allowed
Management Interface Command Set	Configuration Get	Yes
	Configuration Set	
	Controller Health Status Poll	
	Management Endpoint Buffer Read	
	Management Endpoint Buffer Write	
	NVM Subsystem Health Status Poll	
	Read NVMe-MI Data Structure	
	Reset	
	SES Receive	
	SES Send	
	VPD Read	
	VPD Write	
NVMe Admin Command Set 21	Device Self-test	Same restrictions as defined by the NVM Express specification
	Firmware Activate/Commit	
	Firmware Image Download	
	Format NVM	
	Get Features	
	Get Log Page	
	Identify	
	Namespace Attachment	
	Namespace Management	
	Sanitize	
	Security Receive/Send	
	Security Send	
	Set Features	
	Vendor Specific	
	Virtualization Management	
PCIe Command Set	PCIe Configuration Read	Yes
	PCIe Configuration Write	
	PCIe I/O Read	
	PCIe Memory Read	
	PCIe Memory Write	
NOTES:		
<del>1. Refer to the NVM Express specification for the definition of a sanitize operation.</del>		
1. NVMe Admin Commands that are prohibited via the out-of-band mechanism (refer to Figure 110) are not listed since they are always prohibited including during a sanitize operation.		

