



LEGAL NOTICE:

© **Copyright 2008 to 2023 NVM Express®**, Inc. **ALL RIGHTS RESERVED.**

This 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 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: "© 2008 to 2023 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 153rd Drive
Beaverton, OR 97003
USA
info@nvmexpress.org

NVM Express® Technical Proposal for New Feature

Technical Proposal ID	TP4182 Computational Programs Command Set Specification changes to other specifications
Change Date	2023-10-25
Builds on Specification	NVM Express Base Specification 2.0c NVMe-MI Specification 1.2
References Specification	TP4091 Computational Programs

Technical Proposal Author(s)

Name	Company
Kim Malone	Intel
Bill Martin	Samsung

This proposal provides changes to the NVM Express Base Specification and NVMe-MI Specification necessary for the Computational Programs Command Set.

Revision History

Revision Date	Change Description
2023-08-18	<ul style="list-style-type: none">Initial version, moved content from Computational Programs TP4091
2023-08-22	<ul style="list-style-type: none">Added modifications to Command Specific Status ValuesAdded modifications to Command Set Identifiers
2023-10-04	<ul style="list-style-type: none">Addressed review comments from Mike Allison
2023-10-20	<ul style="list-style-type: none">Integrated
2023-10-24	<ul style="list-style-type: none">Removed "I/O" from "Computational Programs Command Set" per Mike Allison
2023-10-25	<ul style="list-style-type: none">Editorial edits integrated per Kim Malone

Summary of changes:

- NVM Express Base Specification
 - Modifications to Theory of Operations types of command sets
 - Modifications to Admin Command Set opcodes
 - Modifications to Command Specific Status Values
 - Modifications to Command Set Identifiers
- NVMe-MI Specification
 - Modifications to Admin Command Set commands
 - Modifications to Log Page Support

Markup Conventions:

Black:	Unchanged (however, hot links are removed)
Blue:	Added
Red Strikethrough:	Deleted
Orange:	Added by TP4171
Blue Highlighted:	TBD values, anchors, and links to be inserted in new text
Green Text:	Editor's notes

Description for Changes Document for NVMe Express Base Specification 2.0c

Modify section 2 as follows:

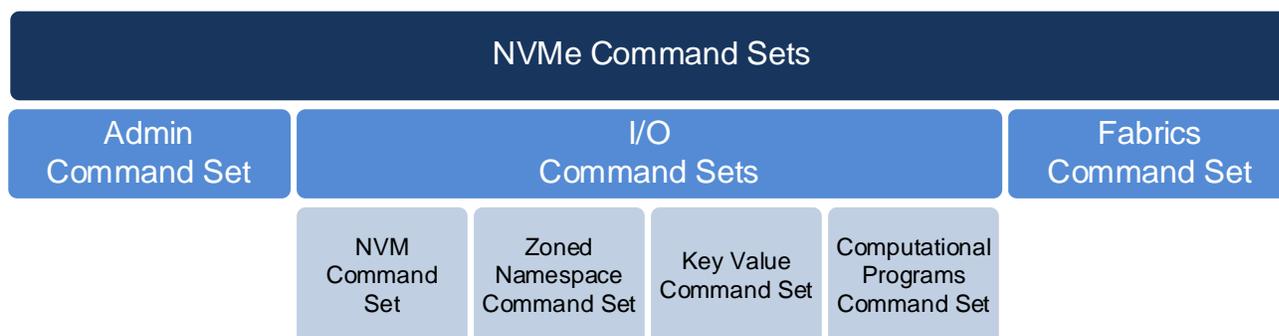
2 Theory of Operations

Editor's note: Changes in this section add the Computational Programs I/O Command Set to lists of command sets.

...

There are three types of commands that are defined in NVMe Express: Admin Commands, I/O Commands, and Fabrics Commands. Figure 5 shows these different command types.

Figure 5: Types of NVMe Command Sets



An Admin Submission Queue and associated Completion Queue exist for the purpose of controller management and control (e.g., creation and deletion of I/O Submission and Completion Queues, aborting commands, etc.). Only commands that are part of the Admin Command Set or the Fabrics Command Set may be submitted to the Admin Submission Queue.

An I/O Command Set is used with an I/O queue pair. This specification defines common I/O commands. I/O Command Sets are defined in NVMe I/O Command Set specifications (e.g., NVM Command Set, Key Value Command Set, ~~or~~ Zoned Namespace Command Set, ~~or~~ Computational Programs Command Set).

...

Modify section 3 as follows:

3 NVMe Express Architecture

...

3.3 NVMe Queue Models

...

3.3.3 Queueing Data Structures

...

3.3.3.2 Common Completion Queue Entry

...

3.3.3.2.1 Status Field Definition

...

3.3.3.2.1.2 Command Specific Status Definition

...

Figure 96: Status Code – Command Specific Status Values, I/O Commands

Value	Description
80h	Conflicting Attributes
81h	Invalid Protection Information
82h	Attempted Write to Read Only Range
83h	Command Size Limit Exceeded
84h	Invalid Command ID
85h	Incompatible Namespace or Format
86h	Fast Copy Not Possible
87h	Overlapping I/O Range
88h	Namespace Not Reachable
89h	Insufficient Resources
8Ah	Insufficient Program Resources
8Bh	Invalid Memory Namespace
8Ch	Invalid Memory Range Set
8Dh	Invalid Memory Range Set Identifier
8Eh	Invalid Program Data
8Fh	Invalid Program Index
90h	Invalid Program Type
91h	Maximum Memory Ranges Exceeded
92h	Maximum Memory Range Sets Exceeded
93h	Maximum Programs Activated
94h	Maximum Program Bytes Exceeded
95h	Memory Range Set In Use
96h	No Program
97h	Overlapping Memory Ranges
98h	Program Not Activated
99h	Program In Use
9Ah	Program Index Not Downloadable
9Bh	Program Too Big
8A9Ch to B7h	Reserved
B8h	Zoned Boundary Error
B9h	Zone Is Full
BAh	Zone Is Read Only
BBh	Zone Is Offline
BCh	Zone Invalid Write
BDh	Too Many Active Zones
BEh	Too Many Open Zones
BFh	Invalid Zone State Transition

...

Modify section 5 as follows:

5 Admin Command Set

Editor's note: Changes to this section add admin commands specific to this I/O Command Set to the opcode table.

Figure 138: Opcodes for Admin Commands

Opcode by Field			Combined Opcode ¹	Namespace Identifier Used ²	Command	Command Set Specific ⁸
(07)	(06:02)	(01:00)				
Generic Command	Function	Data Transfer ³				
0b	000 00b	00b	00h	No	Delete I/O Submission Queue	No
0b	000 00b	01b	01h	No	Create I/O Submission Queue	No
0b	000 00b	10b	02h	Yes	Get Log Page	No
0b	000 01b	00b	04h	No	Delete I/O Completion Queue	No
0b	000 01b	01b	05h	No	Create I/O Completion Queue	No
0b	000 01b	10b	06h	NOTE 6	Identify	No
0b	000 10b	00b	08h	No	Abort	No
0b	000 10b	01b	09h	Yes	Set Features	No
0b	000 10b	10b	0Ah	Yes	Get Features	No
0b	000 11b	00b	0Ch	No	Asynchronous Event Request	No
0b	000 11b	01b	0Dh	Yes	Namespace Management	No
0b	001 00b	00b	10h	No	Firmware Commit	No
0b	001 00b	01b	11h	No	Firmware Image Download	No
0b	001 01b	00b	14h	Yes	Device Self-test	No
0b	001 01b	01b	15h	Yes ⁴	Namespace Attachment	No
0b	001 10b	00b	18h	No	Keep Alive	No
0b	001 10b	01b	19h	Yes ⁵	Directive Send	No
0b	001 10b	10b	1Ah	Yes ⁵	Directive Receive	No
0b	001 11b	00b	1Ch	No	Virtualization Management	No
0b	001 11b	01b	1Dh	No	NVMe-MI Send	No
0b	001 11b	10b	1Eh	No	NVMe-MI Receive	No
0b	010 00b	00b	20h	No	Capacity Management	No
0b	010 01b	00b	24h	No	Lockdown	No

Figure 138: Opcodes for Admin Commands

Opcode by Field			Combined Opcode ¹	Namespace Identifier Used ²	Command	Command Set Specific ⁸
(07)	(06:02)	(01:00)				
Generic Command	Function	Data Transfer ³				
0b	111 11b	00b	7Ch	No	Doorbell Buffer Config	No
0b	111 11b	11b	7Fh	NOTE 9	Fabrics Commands ⁹	No
1b	000 00b	00b	80h	Yes	Format NVM	No
1b	000 00b	01b	81h	NOTE 7	Security Send	No
1b	000 00b	10b	82h	NOTE 7	Security Receive	No
1b	000 01b	00b	84h	No	Sanitize	No
1b	000 01b	01b	85h	Yes	Load Program	CP
1b	000 01b	10b	86h	NOTE 4	Get LBA Status	NVM, ZNS
1b	000 10b	00b	88h	Yes	Program Activation Management	CP
1b	000 10b	01b	89h	Yes	Memory Range Set Management	CP
Vendor Specific						
1b	n/a	NOTE 3	C0h to FFh		Vendor specific	

NOTES:

- Opcodes not listed are reserved.
- A subset of commands use the Namespace Identifier (NSID) field. If the Namespace Identifier field is used, then the value FFFFFFFFh is supported in this field unless otherwise indicated in footnotes in this figure that a specific command does not support that value or supports that value only under specific conditions. When this field is not used, the field is cleared to 0h as described in Figure 87.
- 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 does not support the use of the Namespace Identifier (NSID) field set to FFFFFFFFh.
- Support for the Namespace Identifier field set to FFFFFFFFh depends on the Directive Operation (refer to section 8.7).
- Use of the Namespace Identifier field depends on the CNS value in the Identify Command as described in Figure 273.
- The use of the Namespace Identifier is Security Protocol specific.
- No = Not I/O Command Set specific, A = All I/O Command Sets, NVM = NVM Command Set specific, ZNS = Zoned Namespace Command Set, CP = Computational Programs I/O Command Set.
- All Fabrics commands use the opcode 7Fh. Refer to section 6 for details.

...

5.17 Identify command

5.17.1 Identify command overview

...

Figure 274: Command Set Identifiers

Command Set Identifier Value	Description	Reference Section
00h	NVM Command Set	Refer to the NVM Command Set Specification
01h	Key Value Command Set	Refer to the Key Value Command Set Specification
02h	Zoned Namespace Command Set	Refer to the Zoned Namespace Command Set Specification
03h	Subsystem Local Memory Command Set	Refer to the Subsystem Local Memory Command Set Specification
04h	Computational Programs Command Set	Refer to the Computational Programs Command Set Specification
05h to 2Fh	Reserved	
30h to 3Fh	Vendor specific	
40h to FFh	Reserved	

...

5.17.2.21 Identify I/O Command Set data structure (CNS 1Ch)

...

Figure 290: I/O Command Set Vector

Bit	Description
63:5	Reserved
4	Computational Programs Namespace Command Set: This bit is set to '1' if the Computational Programs Command Set is selected. This bit is cleared to '0' if the Zoned Namespace Command Set is not selected.
3	Subsystem Local Memory Command Set: This bit is set to '1' if the Subsystem Local Memory Command Set is selected. This bit is cleared to '0' if the Subsystem Local Memory Command Set is not selected.
2	Zoned Namespace Command Set: This bit is set to '1' if the Zoned Namespace Command Set is selected. This bit is cleared to '0' if the Zoned Namespace Command Set is not selected.
1	Key Value Command Set: This bit is set to '1' if the Key Value Command Set is selected. This bit is cleared to '0' if the Key Value Command Set is not selected.
0	NVM Command Set: This bit is set to '1' if the NVM Command Set is selected. This bit is cleared to '0' if the NVM Command Set is not selected.

...

Description for Changes Document for NVMe-MI Specification 1.2

Modify section 1 as follows:

1 Introduction

...

1.11 References

...

[NVM Express Computational Programs Command Set Specification, Revision 1.0.](https://www.nvmexpress.org) Available from <https://www.nvmexpress.org>.

...

Modify section 6 as follows:

6 NVM Express Admin Command Set

...

Modify section 6 Figure 116 as follows:

Figure 116: List of NVMe Admin Commands Supported using the Out-of-Band Mechanism

Command	Opcode	NVMe Storage Device O/M/P ¹	NVMe Enclosure O/M/P ¹	Reference Specification
Abort	00h	P	P	NVM Express Base Specification
...
Load Program	85h	P	P	Computational Programs Command Set Specification
Get LBA Status	86h	O	P	NVM Command Set Specification
Program Activation Management	88h	P	P	Computational Programs Command Set Specification
Memory Range Set Management	89h	P	P	Computational Programs Command Set Specification
...

Figure 116: List of NVMe Admin Commands Supported using the Out-of-Band Mechanism

<p>NOTES:</p> <p>1. O/M/P definition: O = Optional, M = Mandatory, P = Prohibited from being supported. An NVMe Enclosure that is also an NVMe Storage Device (i.e., implements Namespaces) shall implement mandatory commands required by either an NVMe Storage Device or an NVMe Enclosure and may implement optional commands allowed by either an NVMe Storage Device or an NVMe Enclosure. Mandatory commands shall be supported using the out-of-band mechanism if the NVMe Controller specified by the Controller ID field supports the command inband.</p> <p>2. If the Retain Asynchronous Event bit is cleared to '0', then the status associated with the NVMe Admin Command shall be Invalid Field in Command (i.e., the NVMe Admin command is aborted). For implementations compliant to version 1.1 or earlier of this specification, the Retain Asynchronous Event bit in the Get Log Page command (refer to the NVM Express Base Specification) may or may not be ignored by the Controller. Refer to section 6.2.</p>

...

6.3 Get Log Page

...

Figure 58: Management Endpoint - Log Page Support

Log Page Name ³	Log Identifier	Support Requirements ¹	
		NVMe Storage Device	NVMe Enclosure
Supported Log Pages	00h	M ²	M ²
Error Information	01h	M	M
SMART / Health Information (Controller scope)	02h	M	O
SMART / Health Information (NVM Subsystem scope)		O	O
Firmware Slot Information	03h	M	O
Changed Namespace List	04h	O	O
Commands Supported and Effects	05h	O	O
Device Self-test	06h	O	O
Telemetry Host-Initiated	07h	O	O
Telemetry Controller-Initiated	08h	O	O
Endurance Group Information	09h	O	O
Predictable Latency Per NVM Set	0Ah	O	O
Predictable Latency Event Aggregate	0Bh	O	O
Asymmetric Namespace Access	0Ch	O	O
Persistent Event	0Dh	O	O
LBA Status Information ⁴	0Eh	O	O
Endurance Group Event Aggregate	0Fh	O	O
Media Unit Status	10h	O	O
Supported Capacity Configuration List	11h	O	O
Feature Identifiers Supported and Effects	12h	M ²	O
NVMe-MI Commands Supported and Effects	13h	O	O
Command and Feature Lockdown	14h	O	O
Boot Partition	15h	O	O
Rotational Media Information	16h	O	O

Figure 58: Management Endpoint - Log Page Support

Log Page Name ³	Log Identifier	Support Requirements ¹	
		NVMe Storage Device	NVMe Enclosure
Discovery	70h	O	O
Reservation Notification	80h	O	O
Sanitize Status	81h	O	O
Program List ⁶	82h	O	O
Downloadable Program Types List ⁶	83h	O	O
Memory Range Set List ⁶	84h	O	O
Changed Zone List ⁵	BFh	O	O
Vendor Specific	C0h to FFh	O	O
Notes: 1. O = Optional, M = Mandatory, P = Prohibited. 2. Optional for versions 1.1 and earlier of this specification. 3. Refer to the NVM Express Base Specification unless another footnote specifies otherwise. 4. Refer to the NVM Command Set Specification. 5. Refer to the Zoned Namespace Command Set Specification. 6. Refer to the Computational Programs Command Set Specification.			

...