



LEGAL NOTICE:

© Copyright 2008 to 2024 NVM Express, Inc. ALL RIGHTS RESERVED.

This erratum 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 erratum, 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 2024 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.
PCI-SIG®, PCI Express®, and PCIe® are registered trademarks of PCI-SIG.
InfiniBand™ is a trademark and servicemark of the InfiniBand Trade Association.

NVM Express Workgroup
c/o VTM Group
3855 SW 153rd 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 © 2008 to 2024 NVM Express, Inc.

NVM Express® Technical Errata

Errata ID	119
Revision Date	01/03/2024
Affected Spec Ver.	NVM Express® Base Specification Revision 2.0c NVM Express® NVM Command Set Specification Revision 1.0c NVM Express® Management Interface Specification Revision 1.2c NVM Express® Key Value Command Set Specification 1.0c
Corrected Spec Ver.	

Errata Author(s)

Name	Company
Mike Allison, Judy Brock, Bill Martin,	Samsung
Fred Knight	NetApp
David Black, Austin Bolen	Dell
Paul Suhler	Kioxia
Andres Baez, Myron Loewen	Solidigm
Randy Jennings	Pure Storage
James Smart	Broadcom

Errata Overview

This ECN updates and clarifies various text within the NVM Express Base Specification Revision 2.0c, NVM Express NVM Command Set Specification Revision 1.0c, NVM Express Management Interface Specification 1.2c, and the NVM Express Key Value Command Set Specification.

Revision History

Revision Date	Change Description
7/19/2023	Initial creation with the solutions for: <ul style="list-style-type: none"> Bug 38 Index offset in the Asymmetric Namespace Access log 0h not defined Bug 138 Capacity Exceeded Status Code Should Be Generic Bug 39 Capacity Management List numbering Bug 22 "Write Protect Until Power Cycle", in a multi-domain system, is this a power cycle of the namespace or the controller?
8/2/2023	Added the solutions for: <ul style="list-style-type: none"> Bug 171 Wording for where subsequent copy data is placed Bug 67 Firmware Commit status code value 12h does not specify what happens when a reset occurs prior to the host re-submitting a commit command Bug 140 Sentence does not parse (grammatically incorrect) Bug 99 Fused Command Processing
8/16/2023	Added the solutions for: <ul style="list-style-type: none"> Bug 89 Reference to Command Specific Information field in Error Information log page points to wrong figure Bug 109 Change "see" to "refer to"
8/23/2023	Added the solutions for: <ul style="list-style-type: none"> Bug 166 Incorrect LBA Status Information log page reference Bug 168 Change "Smart" to "SMART" Bug 165 Mixed cases on Directive Type Bug 192 Async Event Config support for Discovery Controller Bug 147 CC Controller Configuration EN bit - data loss clarification Bug 131 AWUN AWUPF ACWU fields in Base spec missing reference to NVM Command Set spec Bug 118 Controller Reset Clarifications
8/30/2023	Clean up for review in Technical WG.
8/30/2023	Added the solutions for: <ul style="list-style-type: none"> Bug 104 Lockdown Log Contents and Scope Bug 105 Fix section 9.4 (without amendment) Updated bug 138 Capacity Exceeded Status Code Should Be Generic
8/30/2023	Updated bug 104 and bug 105 amendments posted after the Errata Task group meeting. This ECN is ready for review.
9/6/2023	Added the solutions for: <ul style="list-style-type: none"> Bug 183 Allow Stream Resource Allocation Failed to be returned by Streams Directive Enable
9/7/2023	Technical WG review comments.
9/13/2023	As requested by the NVM Express Board of Directors the following bug solution is incorporated: <ul style="list-style-type: none"> Bug 188 Issue with Retain Asynchronous Event bit in NVMe-MI
9/25/2023	Aligned the NVM Express management Interface specification text on the RAE bit to the recommendation in the NVM Express Base Specification that if a log page is not associated to an asynchronous event, then the RAE bit should be cleared to '0'.
10/4/2023	<ul style="list-style-type: none"> fixed grammar in "Description of Changes" for Keep Alive Timeout feature added period to 1.5.TBD Controller Reset entry made ECN117 text orange Removed "For the Controller Reset case of Controller Level Reset," fixed reference to 3.11 Firmware Update Process (was 3.1.1) struck out left over "is ignored" after "controller ignores the UUID index field" corrected color on new figure for Enable Directive.
10/5/2023	<ul style="list-style-type: none"> Firmware Update Process: small fix for blue formatting Stream Resource Allocation Failed: small wording change Reservations: Remove section; no changes
10/11/2023	<ul style="list-style-type: none"> Description of Changes: Stream Resource Allocation Failed as a Clarification. Bring back clarification that command specifying non-zero values of the FUSE field gets aborted.

Technical input submitted to the NVM Express® Workgroup is subject to the terms of the NVM Express® Participant's agreement. Copyright © 2008 to 2024 NVM Express, Inc.

Revision Date	Change Description
11/08/2023	Editorial changes by Errata task group: <ul style="list-style-type: none"> • Updated CSTS.SHST to refer to section 3.6.1. • Remove new completion of Controller Reset text. • Corrected navigation pane to show section hierarchy • Clarified controller enabling after controller shutdown requires CC.SHN to be cleared to 00b. • Added text to RAE bit definition in NVM Express Base specification to refer to NVM Express Management Interface Specification for behaviors per request from Technical Workgroup comment. • Added consistent bullets in the history table,
12/05/2023	Per review feedback from Fred Knight, used the same text pattern for Discovery controllers requirements for commands and features notes as to mandatory and prohibited.
12/22/2023	Integrated
01/03/2023	Editorial updates per Mike Allison and Fred Knight

Description of Changes

NVM Express Base Specification 2.0c:

Backward Incompatible Changes:

- Namespace Write Protect Until Power Cycle is prohibited if the controller and namespace are in different domains because the namespace and controller have to be simultaneously power-cycled to clear that state.

Editorial Changes:

- Clarified the index offset value 0h returning the header of the log page in the Asymmetric Namespace Access log page.
- Clarified the actions on each namespace when an NVM Set is deleted by the Capacity Management command.
- Clarified that the status code 12h does commit a firmware image to the specified slot and that a different Firmware Commit command is needed to activate that image that was committed to the firmware slot.
- Clarified how the controller has a positive indication that completion queues are already being processed.
- Added and corrected references to Error Information Log Entry data structure and Error Information log page.
- Made a consistent use of case for "Directive Type".
- Clarified Keep Alive Timer feature requirements for persistent connections.
- Clarified that Controller Reset does not affect the results of completed commands.
- Added reference to NVM Command Set spec for Identify Controller atomicity fields
- Added correct definition of Controller Reset that includes ability of NVMe-MI to cause a Controller Reset
- Removed incorrect parenthetical CC.EN-based explanations of Controller Reset in favor of definition.
- Added modifications to Controller Level Reset material for completeness and consistency.
- Corrected "Controller Reset" to "Controller Level Reset" for continuing after controller shutdown.
- Clarified descriptions on the UUID index for Vendor Specific Events in the Persistent Event log page.
- Clarified descriptions of what gets read from the Command and Feature Lockdown log page (Log Identifier 14h).
- Clarified Internal Controller Error Handling behavior.
- Clarified that the Stream Resource Allocation Failed status code can be returned if, for a Directive Send Command/Identify Directive/Enable Directive operation, the controller was not able to enable the Streams Directive for the specified namespace and Host Identifier due to insufficient internal resources.
- Use consistent method to describe Discovery Controller requirements associated with explicit persistent connections.

NVM Express NVM Command Set Specification 1.0c:

Editorial Changes:

- Clarified the ordering of copying logical blocks from the different Source Ranges into the destination of the Copy command.
- Clarified the processing by the controller when fused commands are not sequential in the I/O Submission Queue.

Technical input submitted to the NVM Express® Workgroup is subject to the terms of the NVM Express® Participant's agreement. Copyright © 2008 to 2024 NVM Express, Inc.

- Clarified the Compare and Write fused operation by clarifying differences between command and operation.
- Corrected the reference section to the LBA Status Information log page.

NVM Express Management Interface Specification 1.2c:

Backward Incompatible Changes:

- The out-of-band processing of a Get Log Page command is incompatible with a BMC that does not support NVMe-MI 1.2 as that version requires that the RAE bit in the Get Log Page command be set to '1'. BMCs that support NVMe prior to NVMe 1.2, the RAE bit was reserved and those system are required to clear the bit to '0', thus every Get Log Page command is aborted. The solution is to have the Out-of-band processing of a Get Log Page command ignore the RAE bit and process the command as though the RAE bit is set to '1' for log pages used with asynchronous events.

Editorial Changes:

- Updated text to use "refer to" when referring to other specifications, figures, or sections.
- Corrected the reference to the SMART / Health Information log page.
- Corrected the capitalization on "SMART".

NVM Express Key Value Command Set Specification 1.0c:

Editorial Changes:

- Retitled six figures from Command Specific Status Values to Generic Command Status Values, and modified text referring to those figures. All status codes in these tables are generic.

Note:

BLACK text indicates unchanged text. **BLUE** text indicates newly inserted text. **RED stricken** text indicates deleted text; **ORANGE** text indicates changes from another ECN. **Purple** text indicates destination of moved text without changes. **Purple stricken** text indicates source of moved text without changes. **GREEN** text indicates editor notes.

Description of NVM Express Base Specification 2.0c changes

Modify section 1 as shown below:

1 Introduction

1.5 Definitions

...

1.5.TBD Controller Reset

Host modification of the CC property that clears CC.EN from '1' to '0' (refer to section 3.7.2).

...

1.5.15 Directive

A method of information exchange between a host and either an NVM subsystem or a controller. Information may be transmitted using the Directive Send and Directive Receive commands. A subset of I/O commands may include a Directive Type field and a Directive Specific field to communicate more information that is specific to the associated I/O command. Refer to section 8.7.

...

Modify section 3 as shown below:

3 Admin Command Set

3.1 NVM Controller Architecture

...

3.1.2 Controller Types

...

3.1.2.3 Discovery Controller

...

3.1.2.3.2 Command Support

...

Figure 32: Discovery Controller – Admin Command Support

Command	Command Support Requirements ¹	Reference
...		
Set Features	NOTE 2	5.27
Get Features	NOTE 2	5.15
Asynchronous Event Request	NOTE 2	5.2
...		
Keep Alive	NOTE 2	5.18
...		
Notes:		
1. O/M/P definition: O = Optional, M = Mandatory, P = Prohibited		

Technical input submitted to the NVM Express® Workgroup is subject to the terms of the NVM Express® Participant's agreement. Copyright © 2008 to 2024 NVM Express, Inc.

Figure 32: Discovery Controller – Admin Command Support

Command	Command Support Requirements ¹	Reference
2.	For Discovery controllers that support explicit persistent connections, this command is mandatory. For Discovery controllers that do not support explicit persistent connections, the this command is prohibited. For Discovery controllers that support explicit persistent connections, the command is mandatory.	
3.	Mandatory for CDCs and optional for Discovery controllers that are not a CDC.	

3.1.2.3.3 Features Support

...

Figure 34: Discovery Controller – Feature Support

Feature Name	Feature Support Requirements ¹	Logged in Persistent Event Log ¹
...		
Asynchronous Event Configuration	NOTE TBD	NR
...		
Keep Alive Timer	M TBD NOTE TBD	O
...		
Notes:		
<p>...</p> <p>TBD. For Discovery controllers that support explicit persistent connections, this feature is mandatory. For Discovery controllers that do not support explicit persistent connections, the Keep Alive Timer this feature is prohibited.</p> <p><Orange text is from ECN117></p>		

...

3.1.3.5 Offset 14h: CC – Controller Configuration

...

Figure 46: Offset 14h: CC – Controller Configuration

Bits	Type	Reset	Description
...			
00	RW	0b	<p>Enable (EN): When set to '1', then the controller shall process commands. When cleared to '0', then the controller shall not process commands nor post completion queue entries to Completion Queues. When the host modifies CC to clear this bit transitions from '1' to '0', the controller is reset (i.e., a Controller Reset, refer to section 3.7.2). That reset deletes all I/O Submission Queues and I/O Completion Queues, resets the Admin Submission Queue and Completion Queue, and brings the hardware to an idle state. That reset does not affect transport specific state (e.g. PCI Express registers including MMIO MSI-X registers), nor the Admin Queue properties (AQA, ASQ, or ACQ). All other controller properties defined in this section and internal controller state (e.g., Feature values defined in section 5.27.1 that are not persistent across power states) are reset to their default values. The controller shall ensure that there is no impact (e.g., data loss) for caused by that Controller Reset to the results of commands that have had corresponding completion queue entries posted to an I/O Completion Queue prior to that Controller Reset. Refer to section 3.6.</p> <p>...</p>

...

Figure 47: Offset 1Ch: CSTS – Controller Status

Bits	Type	Reset ¹	Description										
03:02	RO	00b	Shutdown Status (SHST): This field indicates the status of shutdown processing that is initiated by the host setting the CC.SHN field, the host setting the NSSC property, or a Management Endpoint has processed an NVMe-MI Shutdown command (refer to the NVM Express Management Interface specification).										
			The shutdown status values are defined as:										
			<table><tr><th>Value</th><th>Definition</th></tr><tr><td>00b</td><td>Normal operation (no shutdown has been requested)</td></tr><tr><td>01b</td><td>Shutdown processing occurring</td></tr><tr><td>10b</td><td>Shutdown processing complete</td></tr><tr><td>11b</td><td>Reserved</td></tr></table>	Value	Definition	00b	Normal operation (no shutdown has been requested)	01b	Shutdown processing occurring	10b	Shutdown processing complete	11b	Reserved
			Value	Definition									
			00b	Normal operation (no shutdown has been requested)									
			01b	Shutdown processing occurring									
			10b	Shutdown processing complete									
			11b	Reserved									
			If this field is set to 01b, then:										
			<ul style="list-style-type: none">an NVM Subsystem Reset aborts a controller shutdown and an NVM Subsystem Shutdown; andany other type of Controller Level Reset:<ul style="list-style-type: none">may or may not abort a controller shutdown; andshall not abort an NVM Subsystem Shutdown.										
If this field is set to 01b when a Controller Level Reset is initiated and the shutdown is not aborted, then this field transitions to 00b on the reset and then to 01b to indicate the shutdown is still in progress and host software may or may not observe this transition.													
If CSTS.ST is cleared to '0' and this field is set to 10b, then to start executing commands on the controller :													
<ul style="list-style-type: none">If CC.EN is set to '1', to start executing commands on the controller after a shutdown operation (CSTS.SHST set to 10b), a Controller Level Reset (e.g., a Controller Reset) (CC.EN cleared to '0') is required. If host software submits commands to the controller without issuing a prior Controller Level Reset, the behavior is undefined; andIf CC.EN is cleared to '0', to start executing commands on the controller:<ul style="list-style-type: none">a Controller Level Reset is required; orCC.EN is required to be set to '1' and CC.SHN is required to be cleared to 00b with the same write to the CC property (refer to section 3.6.1).													
If CSTS.ST is set to '1' and this field is set to 10b, then an NVM Subsystem Reset is required to start executing commands.													
Refer to section 3.6.3 on the reset behavior of this field when CAP.CPS is set to 10b or 11b.													

...

3.3 NVM Queue Models

...

3.3.2 Message-based Transport Queue Model

...

3.3.2.9 Transport Requirements

...

Technical input submitted to the NVM Express® Workgroup is subject to the terms of the NVM Express® Participant's agreement. Copyright © 2008 to 2024 NVM Express, Inc.

The NVMe Transport may support NVMe Transport error detection and report errors to the NVMe layer in command status values. The controller may record NVMe Transport specific errors in the Error Information log page (refer to section 5.16.1.2). Transport errors that cause loss of a message or loss of data in a way that the low-level NVMe Transport cannot replay or recover should cause:

- the deletion of the individual I/O Queues (refer to section 3.3.2.4) and the associated NVMe Transport connection on which that NVMe Transport level error occurred; or
- termination of the NVMe Transport connection and the association between the host and controller.

...

3.4 Command Architecture Submission and Completion Mechanism

...

3.4.2 Fused Operations

Fused operations enable a more complex command by “fusing” together two simpler commands. This feature is optional; support for this feature is indicated in the FUSES field in the Identify Controller data structure in Figure 275.

Whether a command is part of a fused operation is indicated in specified by the Fused Operation (FUSE) field of Command Dword 0 shown in Figure 86. The FUSE field also indicates specifies whether each the command is the first command in the fused operation or the second command in the fused operation. If the FUSE field is set to a non-zero value and the controller does not support the requested fused operation, then the controller should abort the command with a status code of Invalid Field in Command.

In a fused operation, the requirements are:

- The commands shall be executed in sequence as an atomic unit. The controller shall behave as if no other operations have been executed between these two commands;
- The operation ends at the point an error is encountered in either command. If the first command in the sequence failed, then the second command in the sequence shall be aborted. If the second command in the sequence failed, then the completion status of the first command is sequence specific and is defined within the Fused Operation section of the applicable I/O Command Set specification;
- The commands shall be inserted next to each other in the same Submission Queue. If the controller processes a command violating this condition (e.g., a command with the FUSE field cleared to 00b (i.e., Normal operation) is inserted immediately after a command specifying the FUSE field set to 01b (i.e., Fused operation, first command), or is inserted immediately before a command specifying the FUSE field set to 10b (i.e., Fused operation, second command)), then the controller shall abort the command specifying non-zero values of the FUSE field with a status code of Command Aborted due to Missing Fused Command. If the first command is in the last slot in the Submission Queue, then the second command shall be in the first slot in the Submission Queue as part of wrapping around. In the memory-based transport queue model, the Submission Queue Tail doorbell pointer update shall indicate both commands as part of one doorbell update. In the message-based transport queue model, the command capsules shall be submitted in-order.
- To abort the fused operation, the host shall submit submits an Abort command separately for each of the commands; and
- A completion queue entry is posted by the controller for each of the commands.

~~Whether a command is part of a fused operation is indicated in the Fused Operation (FUSE) field of Command Dword 0 shown in Figure 86. The FUSE field also indicates whether each command is the first command in the fused operation or the second command in the fused operation. If the FUSE field is set to a non-zero value and the controller does not support the requested fused operation, then the controller should abort the command with a status code of Invalid Field in Command.~~

Refer to each I/O Command Set specification for applicability and additional details, if any.

...

3.6 Shutdown Processing

...

3.6.1 Memory-based Transport Controller Shutdown

...

To start executing commands on the controller after that controller reports controller shutdown processing complete (i.e., CSTS.ST is cleared to '0' and CSTS.SHST is set to 10b) utilizing CC.EN:

- if CC.EN is set to '1', then a Controller Level Reset (~~CC.EN cleared from '1' to '0'~~) is required to ~~clear CC.EN to '0'~~ on that controller;~~or~~ and
- if CC.EN is cleared to '0', then the controller is required to be enabled (i.e., CC.EN is set to '1' from '0'). ~~The CC.SHN is required to be cleared to 00b in the same write to the CC property.~~

The initialization sequence should then be executed on that controller.

...

3.7 Resets

...

3.7.2 Controller Level Reset

The following methods initiate a Controller Level Reset:

- NVM Subsystem Reset;
- Controller Reset (i.e., ~~CC.EN~~ host clears ~~transitions~~ CC.EN from '1' to '0'); and
- Transport specific reset types (refer to the applicable NVMe Transport binding specification), if any.

A Controller Level Reset consists of the following actions:

- The controller stops processing any outstanding Admin or I/O commands;
 - All I/O Submission Queues are deleted;
 - All I/O Completion Queues are deleted;
 - The controller is brought to an ~~idle~~ idle state. When this is complete, CSTS.RDY is cleared to '0'; and
 - All controller properties defined in section 3.1.3 and internal controller state are reset, with the following exceptions:
 - for controllers using a memory-based transport:
 - the Admin Queue properties (AQA, ASQ, or ACQ) are not reset as part of a Controller Reset;
 - the Controller Memory Buffer Memory Space Control property (CMBMSC) is reset as part of neither a Controller Reset nor a Function Level Reset; and
 - the Persistent Memory Region Memory Space Control Upper property (PMRMSCU) and the Persistent Memory Region Memory Space Control Lower property (PMRMSCL) are not reset as part of a Controller Reset;
- and
- for controllers using a message-based transport:
 - there are no exceptions.

In all Controller Level Reset cases except a Controller Reset, the controller properties are reset as defined by the applicable NVMe Transport binding specification.

To continue after a Controller Level Reset, the host ~~shall~~ should:

Technical input submitted to the NVM Express® Workgroup is subject to the terms of the NVM Express® Participant's agreement. Copyright © 2008 to 2024 NVM Express, Inc.

- **U**update transport specific state and controller property state as appropriate;
- **S**et CC.EN to '1';
- **W**ait for CSTS.RDY to be set to '1';
- **C**onfigure the controller using Admin commands as needed;
- **C**reate I/O Completion Queues and I/O Submission Queues as needed; and
- **P**roceed with normal I/O operations.

Note that all Controller Level Reset cases except a Controller Reset result in the controller immediately losing communication with the host. In all these cases, the controller is unable to indicate any aborts or update any completion queue entries.

...

3.11 Firmware Update Process

...

3. The last step is to perform a reset that then causes the firmware image specified in the Firmware Slot field in the Firmware Commit command to be activated. The reset may be an NVM Subsystem Reset, Conventional Reset, Function Level Reset, or Controller Reset (~~CC.EN transitions from '1' to '0'~~):
 - a. In some cases a Conventional Reset or NVM Subsystem Reset is required to activate a firmware image. This requirement is indicated by Firmware Commit command specific status (refer to section 5.12.1);

...

The process for a firmware update to be activated on a domain without a reset is:

1. The host issues a Firmware Image Download command to download the firmware image to a controller. There may be multiple portions of the firmware image to download, thus the offset for each portion of the firmware image being downloaded on that controller is specified in the Firmware Image Download command. The data provided in the Firmware Image Download command should conform to the Firmware Update Granularity indicated in the Identify Controller data structure or the firmware update may fail;
2. The host submits a Firmware Commit command on that controller with a Commit Action of 011b which specifies that the firmware image should be activated immediately without reset. The downloaded firmware image should replace the firmware image in the firmware slot. If no firmware image was downloaded since the last reset or Firmware Commit command, (i.e., the first step was skipped), then that controller shall verify and activate the firmware image in the specified slot. If that controller starts to activate the firmware image, any controllers affected by the new firmware image send a Firmware Activation Starting asynchronous event to the host if Firmware Activation Notices are enabled (refer to Figure 326):
 - a. The Firmware Commit command may also be used to activate a firmware image associated with a previously committed firmware slot;
3. The controller completes the Firmware Commit command. The following actions are taken in certain error scenarios:
 - a. If the firmware image is invalid, then the controller **aborts the command with an reports the appropriate status code error** (e.g., Invalid Firmware Image);
 - b. If the firmware activation was not successful because a Controller Level Reset is required to activate this firmware **image**, then the controller **aborts the command with a status code reports an-error** of Firmware Activation Requires Controller Level Reset and the firmware image is applied at the next Controller Level Reset;
 - c. If the firmware activation was not successful because an NVM Subsystem Reset is required to activate this firmware image, then the controller **aborts the command with a status code reports an-error** of Firmware Activation Requires NVM Subsystem Reset and the firmware image is applied at the next NVM Subsystem Reset;
 - d. If the firmware activation was not successful because a Conventional Reset is required to activate this firmware **image**, then the controller **aborts the command with a status code reports**

Technical input submitted to the NVM Express® Workgroup is subject to the terms of the NVM Express® Participant's agreement. Copyright © 2008 to 2024 NVM Express, Inc.

~~an error~~ of Firmware Activation Requires Conventional Reset and the firmware image is applied at the next Conventional Reset; and

- e. If the firmware activation was not successful because the firmware activation time would exceed the MTFA value reported in the Identify Controller data structure, then the controller ~~aborts the command with a status code reports an error~~ of Firmware Activation Requires Maximum Time Violation. In this case, ~~to activate the firmware, the Firmware Commit command needs to be re-issued and the firmware image activated using a reset~~ the firmware image was committed to the specified firmware slot. To activate that firmware image, the host may issue a Firmware Commit command that specifies:

- i. a Commit Action set to 010b (i.e., activate using a Controller Level Reset); and
- ii. the same firmware slot.

...

Modify section 5 as shown below:

5 Admin Command Set

...

5.3 Capacity Management command

...

5.3.2 Endurance Group Operations

...

If the Operation field specifies the Create Endurance Group operation and the Capacity Lower and Capacity Upper fields specify a value that requires allocation of NVM capacity that is greater than the value in:

- a) the Max Endurance Group Capacity (MEGCAP) field in the Identify Controller data structure;
- b) the Unallocated NVM Capacity (UNVMCAP) field in the Identify Controller data structure; or
- c) the Max Endurance Group Capacity (MEGCAP) field in the Domain Attributes Entry for the domain in which the Endurance Group is being created,

then the controller:

- a) shall abort the command with Insufficient Capacity status; and
- b) if the Error Information log page is supported, shall indicate in the Command Specific Information field in the Error Information Log Entry data structure (refer to Figure 206) the total amount of NVM capacity in bytes of the largest Endurance Group that is able to be created.

...

5.3.3 NVM Set Operations

...

If the Operation field specifies the Create NVM Set operation and the Capacity Lower and Capacity Upper fields specify a value that requires allocation of NVM capacity that is greater than the value in the Unallocated Endurance Group Capacity (UEGCAP) field in the Endurance Group Information log page (refer to Figure 217) for the specified Endurance Group, then the controller:

- a) if the Error Information log page is supported, shall indicate in the Command Specific Information field in the Error Information Log Entry data structure (refer to Figure 206) the total amount of NVM capacity in bytes of the largest NVM Set that is able to be created; and
- b) shall abort the command with Insufficient Capacity status.

...

Technical input submitted to the NVM Express® Workgroup is subject to the terms of the NVM Express® Participant's agreement. Copyright © 2008 to 2024 NVM Express, Inc.

5.3.4 Command Completion

Upon completion of the Capacity Management command, the controller posts a completion queue entry to the Admin Completion Queue. Capacity Management command specific status values are defined in Figure 153.

Figure 153: Capacity Management – Command Specific Status Values

Value	Description
26h	Insufficient Capacity: The requested operation requires more free space than is currently available. The Command Specific Information field in the Error Information Log Entry data structure (refer to Figure 206) of the Error Information log page (refer to Figure 193) specifies the total amount of NVM capacity in bytes required to create the Endurance Group or NVM Set.
2Dh	Identifier Unavailable: The number of Endurance Groups or NVM Sets supported has been exceeded.

...

5.12 Firmware Commit command

...

5.12.1 Command Completion

...

Figure 183: Firmware Commit – Command Specific Status Values

Value	Description
...	
11h	Firmware Activation Requires Controller Level Reset: The firmware commit was successful; however, the firmware image specified does not support being activated without a Controller Level Reset. The firmware image shall be activated at the next Controller Level Reset. This status code should be returned only if the Commit Action field in the Firmware Commit command is set to 011b (i.e., activate immediately).
12h	Firmware Activation Requires Maximum Time Violation: The image specified if activated immediately would exceed the Maximum Time for Firmware Activation (MTFA) value reported in the Identify Controller data structure (refer to Figure 275). To activate the firmware, the Firmware Commit command needs to be re-issued and the image activated using a reset. The firmware commit was successful; however, the firmware image was not activated. Activation of the specified firmware image requires more time than the amount indicated in the Maximum Time for Firmware Activation (MTFA) field in the Identify Controller data structure (refer to Figure 275). To activate the firmware image committed to the specified firmware slot, a Firmware Commit command specifying a Commit Action set to 010b may be issued as described in section 3.11.
...	

...

5.16 Get Log Page command

...

Figure 197: Get Log Page – Command Dword 10

Bits	Description
31:16	Number of Dwords Lower (NUMDL): This field specifies the least significant 16 bits of the number of dwords to return unless otherwise specified. If host software specifies a size larger than the log page requested, the controller returns the complete log page with undefined results for dwords beyond the end of the log page. The combined NUMDL and NUMDU fields form a 0's based value.

Figure 197: Get Log Page – Command Dword 10

Bits	Description
15	<p>Retain Asynchronous Event (RAE): This bit specifies whether to retain or clear an Asynchronous Event. If this bit is cleared to '0', the corresponding Asynchronous Event is cleared by the controller upon successful command completions. If this bit is set to '1', the corresponding Asynchronous Event is retained (i.e., not cleared) by the controller upon command completion.</p> <p>If the command does not complete successfully, the Asynchronous Event shall be retained by the controller.</p> <p>Host software should clear this bit to '0' for log pages that are not used with Asynchronous Events. Refer to section 5.2.</p> <p>Refer to the NVM Express Admin Command Set section of the NVM Express Management Interface Specification for the behavior associated with this bit when the Get Log Page command is received on a Management Endpoint.</p>
14:08	Log Specific Parameter (LSP): If not defined for the log specified by the Log Page Identifier field, this field is reserved.
07:00	Log Page Identifier (LID): This field specifies the identifier of the log page to retrieve.

...

5.16.1 Log Specific Information

...

5.16.1.9 Telemetry Controller-Initiated (Log Identifier 08h)

This log consists of a header describing the log and zero or more Telemetry Data Blocks (refer to section 8.23). All Telemetry Data Blocks are 512 bytes in size. This log is a controller initiated capture of the controller's internal state. The Telemetry Controller-Initiated Data for Data Area 1 through Data Area 3 shall persist across all resets. The Telemetry Controller-Initiated Data for Data Area 4 may persist across [Controller Level Resets](#) ~~controller resets~~. If the host specifies a Log Page Offset Lower value that is not a multiple of 512 bytes in the Get Log Page command for this log, then the controller shall return an error of Invalid Field in Command. This log page is global to the controller.

...

5.16.1.13 Asymmetric Namespace Access (Log Identifier 0Ch)

This log consists of a header describing the log and descriptors containing the asymmetric namespace access information for ANA Groups (refer to section 8.1.2) that contain namespaces that are attached to the controller processing the command. If ANA Reporting (refer to section 8.1) is supported, this log page is supported. ANA Group Descriptors shall be returned in ascending ANA Group Identifier order.

If the Index Offset Supported bit is cleared to '0' in the LID Support and Effects data structure for this log page (refer to Figure 204), then:

- if the RGO bit is cleared to '0' in Command Dword 10, then the LPOL field in Command Dword 12 and the LPOU field in Command Dword 13 of the Get Log Page command should be cleared to 0h.

~~If the Index Offset Supported bit is set to '1' in the LID Supported and Effects data structure for this log page (refer to Figure 204), then:~~

- ~~the entry data structure that is indexed is an ANA Group Descriptor (e.g., specifying an index offset of 2 returns this log page starting at the offset of ANA Group Descriptor 1).~~

~~If the Index Offset Supported bit is set to '1' in the LID Supported and Effects data structure for this log page (refer to Figure 204) and the OT bit is set to '1', then, an index value specified in the Log Page Offset Lower (LPOL) field (refer to [Figure 224](#)) and Log Page Offset Upper (LPOU) field (refer to [Figure 225](#)) in the Get Log Page command shall be used to index to the header or an ANA Group Descriptor within the list of ANA~~

Technical input submitted to the NVM Express® Workgroup is subject to the terms of the NVM Express® Participant's agreement. Copyright © 2008 to 2024 NVM Express, Inc.

Group Descriptors as shown in Figure 221 (e.g., specifying an index offset of 0 returns this log page starting at byte offset 0h, specifying an index offset of 1 returns this log page starting at the offset of ANA Group Descriptor 0, specifying an index offset of 2 returns this log page starting at the offset of ANA Group Descriptor 1, etc.).

If the host performs multiple Get Log Page commands to read the ANA log page (e.g., using the LPOL field or the LPOU field), the host should re-read the header of the log page and ensure that the Change Count field in the Asymmetric Namespace Access log matches the original value read. If it does not match, then the data captured is not consistent and the ANA log page should be re-read.

The Log Specific Parameter field in Command Dword 10 (refer to Figure 197) for this log page is defined in Figure 220.

...

Figure 221: Asymmetric Namespace Access Log Page

Bytes	Description
Header	
07:00	Change Count: This field contains a 64-bit incrementing Asymmetric Namespace Access log change count, indicating an identifier for this set of asymmetric namespace access information. The count starts at 0h following a Controller Level Reset and is incremented each time the contents of the log page change (e.g., not only if an Asymmetric Namespace Access Change Asynchronous Event Notification is generated). If the value of this field is FFFFFFFF_FFFFFFFFh, then the field shall be cleared to 0h when incremented (i.e., rolls over to 0h).
09:08	Number of ANA Group Descriptors (NAGD): This field indicates the number of ANA Group Descriptors available in the log page. The log page shall contain one ANA Group Descriptor for each ANA Group that contains namespaces that are attached to the controller. If, for an ANA Group, there are no namespaces attached to the controller processing the command, then no ANA Group Descriptor is returned for that ANA Group (i.e., an ANA Group Descriptor is returned only if that ANA Group contains namespaces that are attached to the controller processing the command). If no namespaces are attached to the controller, then the log page does not contain any ANA Group Descriptors and this field is cleared to 0h.
15:10	Reserved
List of ANA Group Descriptors	
n:16	ANA Group Descriptor 0, if any
m:n+1	ANA Group Descriptor 1, if any
...	...
x:y	ANA Group Descriptor NAGD-1n, if any

...

5.16.1.14 Persistent Event (Log Identifier 0Dh)

...

5.16.1.14.1 Persistent Event Log Events

...

5.16.1.14.1.14 Vendor Specific Event (Event Type DEh)

...

If a UUID Index is specified in the Get Log Page command (refer to section 5.16), then the controller shall return:

- a) Vendor specific events defined by the vendor identified by the specified based on that UUID index; and

Technical input submitted to the NVM Express® Workgroup is subject to the terms of the NVM Express® Participant's agreement. Copyright © 2008 to 2024 NVM Express, Inc.

- b) Vendor specific events defined by the NVM subsystem manufacturer.

...

5.16.1.20 Command and Feature Lockdown (Log Identifier 14h)

...

Figure 259: Command and Feature Lockdown Log Specific Parameter Field

Bits	Description														
14	Reserved														
13:12	<p>Contents (CNTTS): This field in combination with the Scope field specifies the contents of the Command and Feature Identifier List field in the log page.</p> <table> <tr> <th>Value</th><th>Command and Feature Identifier List Definition</th></tr> <tr> <td>00b</td><td>List of command opcodes or Feature Identifiers based on specified by the Scope field that are supported able to be prohibited.</td></tr> <tr> <td>01b</td><td>List of command opcodes or Feature Identifiers based on specified by the Scope field that are currently prohibited if received on an NVM Express controller Admin submission queue.</td></tr> <tr> <td>10b</td><td>List of command opcodes or Feature Identifiers based on specified by the Scope field that are currently prohibited if received out-of-band on a Management Endpoint.</td></tr> <tr> <td>11b</td><td>Reserved</td></tr> </table>	Value	Command and Feature Identifier List Definition	00b	List of command opcodes or Feature Identifiers based on specified by the Scope field that are supported able to be prohibited.	01b	List of command opcodes or Feature Identifiers based on specified by the Scope field that are currently prohibited if received on an NVM Express controller Admin submission queue.	10b	List of command opcodes or Feature Identifiers based on specified by the Scope field that are currently prohibited if received out-of-band on a Management Endpoint.	11b	Reserved				
Value	Command and Feature Identifier List Definition														
00b	List of command opcodes or Feature Identifiers based on specified by the Scope field that are supported able to be prohibited.														
01b	List of command opcodes or Feature Identifiers based on specified by the Scope field that are currently prohibited if received on an NVM Express controller Admin submission queue.														
10b	List of command opcodes or Feature Identifiers based on specified by the Scope field that are currently prohibited if received out-of-band on a Management Endpoint.														
11b	Reserved														
11:08	<p>Scope (SCP): This field in combination with the Contents field specifies the contents of the Command and Feature Identifier List field in the log page.</p> <table> <tr> <th>Value</th><th>Command and Feature Identifier List Contents</th></tr> <tr> <td>0h</td><td>List of Admin Command Set opcodes</td></tr> <tr> <td>1h</td><td>Reserved</td></tr> <tr> <td>2h</td><td>List of Feature Identifiers</td></tr> <tr> <td>3h</td><td>List of a Management Interface Command Set opcodes (refer to the NVM Express Management Interface Specification)</td></tr> <tr> <td>4h</td><td>List of a PCIe Command Set opcodes (refer to the NVM Express Management Interface Specification)</td></tr> <tr> <td>5h to Fh</td><td>Reserved</td></tr> </table>	Value	Command and Feature Identifier List Contents	0h	List of Admin Command Set opcodes	1h	Reserved	2h	List of Feature Identifiers	3h	List of a Management Interface Command Set opcodes (refer to the NVM Express Management Interface Specification)	4h	List of a PCIe Command Set opcodes (refer to the NVM Express Management Interface Specification)	5h to Fh	Reserved
Value	Command and Feature Identifier List Contents														
0h	List of Admin Command Set opcodes														
1h	Reserved														
2h	List of Feature Identifiers														
3h	List of a Management Interface Command Set opcodes (refer to the NVM Express Management Interface Specification)														
4h	List of a PCIe Command Set opcodes (refer to the NVM Express Management Interface Specification)														
5h to Fh	Reserved														

If a UUID Index is specified in the Get Log Page command (refer to section 5.16) with the Scope field is set to 2h, then the controller should return vendor specific Set Features lockdown information ~~defined by the vendor identified by the specified~~ based on that UUID index ~~field~~. If the Scope field is not set to 2h, then the controller ignores the UUID index field ~~is ignored~~.

...

5.17 Identify command

...

5.17.2 Identify Data Structures

...

5.17.2.1 Identify Controller Data Structure (CNS 01h)

...

Technical input submitted to the NVM Express® Workgroup is subject to the terms of the NVM Express® Participant's agreement. Copyright © 2008 to 2024 NVM Express, Inc.

Figure 275: Identify – Identify Controller Data Structure, I/O Command Set Independent

Bytes	I/O ¹	Admin ¹	Disc ¹	Description
...				
527:526	M	R	R	Atomic Write Unit Normal (AWUN): This field is specific to namespaces that are associated with command sets that specify logical blocks (i.e., Command Set Identifier 0h or 2h), and shall be cleared to 0h for namespaces that are not associated with command sets that specify logical blocks. Refer to the applicable I/O Command Set specification (e.g., the Atomic Operation section of the NVM Command Set specification).
529:528	M	M	R	Atomic Write Unit Power Fail (AWUPF): This field is specific to namespaces that are associated with command sets that specify logical blocks (i.e., Command Set Identifier 0h or 2h), and shall be cleared to 0h for namespaces that are not associated with command sets that specify logical blocks. Refer to the applicable I/O Command Set specification (e.g., the Atomic Operation section of the NVM Command Set specification).
...				
533:532	O	R	R	Atomic Compare & Write Unit (ACWU): This field is specific to namespaces that are associated with command sets that specify logical blocks (i.e., Command Set Identifier 0h or 2h), and shall be cleared to 0h for namespaces that are not associated with command sets that specify logical blocks. Refer to the applicable I/O Command Set specification (e.g., the Atomic Operation section of the NVM Command Set specification).
...				

...

5.17.2.13 Primary Controller Capabilities data structure (CNS 14h)

...

Figure 281: Identify – Primary Controller Capabilities Structure

Bytes	Description
...	...
41:40	VQ Resources Flexible Allocated to Primary (VQRFAP): This field indicates the total number of VQ Flexible Resources currently allocated to the primary controller. This value may change after a Controller Level Reset other than a Controller Reset (i.e., CC.EN transitions from '1' to '0') if a new value was set using the Virtualization Management command. The default value of this field is implementation specific.
...	...
73:72	VI Resources Flexible Allocated to Primary (VIRFAP): This field indicates the total number of VI Flexible Resources currently allocated to the primary controller. This value may change after a Controller Level Reset other than a Controller Reset (i.e., CC.EN transitions from '1' to '0') if a new value was set using the Virtualization Management command. The default value of this field is implementation specific.
...	...

...

5.23 Namespace Management command

...

5.23.1 Command Completion

...

Technical input submitted to the NVM Express® Workgroup is subject to the terms of the NVM Express® Participant's agreement. Copyright © 2008 to 2024 NVM Express, Inc.

Figure 301: Namespace Management – Command Specific Status Values

Value	Description
...	
15h	Namespace Insufficient Capacity: Creating the namespace requires more unallocated capacity than is currently available. The Command Specific Information field in the Error Information Log Entry data structure (refer to Figure 206) of the Error Information log page specifies the total amount of unallocated NVM capacity required to create the namespace in bytes.
...	

...

5.27 Set Features command

...

5.27.1 Feature Specific Information

...

5.27.1.6 Interrupt Coalescing (Feature Identifier 08h)

...

If the controller detects that interrupts are already being processed for this vector, then the controller may delay additional interrupts. Specifically, if the Completion Queue Head Doorbell property ~~is being updated~~ that is associated with a particular interrupt vector ~~is being updated~~, then the controller has a positive indication that completion queue entries are already being processed. In this case, the aggregation time and/or the aggregation threshold may be reset/restarted upon the associated property write. This may result in interrupts being delayed indefinitely in certain workloads where the aggregation time or aggregation threshold is non-zero.

...

5.27.1.28 Namespace Write Protection Config (Feature Identifier 84h)

...

If a Set Features command attempts to change the namespace write protection state of a namespace that is in the Write Protect Until Power Cycle state or the Permanent Write Protect state, then the ~~command~~ controller shall abort the command with a status code of Feature Not Changeable.

If a Set Features command attempts to change the namespace write protection state of a namespace to the Write Protect Until Power Cycle state and bit 0 of the of the Write Protection Authentication Control field is cleared to '0', then the ~~command~~ controller shall abort the command with a status code of Feature Not Changeable.

If a Set Features command attempts to change the namespace write protection state of a namespace to the Write Protect Until Power Cycle state in a multi-domain NVM subsystem (i.e., the MDS bit is set to '1' in the CTRATT field of the Identify Controller data structure (refer to Figure 275)), then the controller should abort the command with a status code of Feature Not Changeable.

<Note to Editor: Change above “should” to “shall” in next numbered (non-lettered) release>:

If a Set Features command changes the namespace to a write protected state, then the controller shall commit all volatile write cache data and metadata associated with the specified namespace to non-volatile media as part of transitioning to the write protected state.

...

5.28 Virtualization Management command

...

Technical input submitted to the NVM Express® Workgroup is subject to the terms of the NVM Express® Participant's agreement. Copyright © 2008 to 2024 NVM Express, Inc.

Figure 371: Virtualization Management – Command Dword 10

Bits	Description								
...	...								
03:00	<p>Action (ACT): This field indicates the operation for the command to perform as described below.</p> <table> <tr> <th>Value</th><th>Description</th></tr> <tr> <td>0h</td><td>Reserved</td></tr> <tr> <td>1h</td><td>Primary Controller Flexible Allocation: Set the number of Flexible Resources allocated to this primary controller following the next Controller Level Reset other than a Controller Reset (i.e., CC-EN transitions from '1' to '0'). If the Controller Identifier field does not correspond to this primary controller, then a status code of Invalid Controller Identifier is returned. This value is persistent across power cycles and resets.</td></tr> <tr> <td>...</td><td>...</td></tr> </table>	Value	Description	0h	Reserved	1h	Primary Controller Flexible Allocation: Set the number of Flexible Resources allocated to this primary controller following the next Controller Level Reset other than a Controller Reset (i.e., CC-EN transitions from '1' to '0'). If the Controller Identifier field does not correspond to this primary controller, then a status code of Invalid Controller Identifier is returned. This value is persistent across power cycles and resets.
Value	Description								
0h	Reserved								
1h	Primary Controller Flexible Allocation: Set the number of Flexible Resources allocated to this primary controller following the next Controller Level Reset other than a Controller Reset (i.e., CC-EN transitions from '1' to '0'). If the Controller Identifier field does not correspond to this primary controller, then a status code of Invalid Controller Identifier is returned. This value is persistent across power cycles and resets.								
...	...								

...

Modify section 6 as shown below:

6 Fabrics Command Set

...

6.3 Connect Command and Response

...

For a Connect command that fails, the controller shall not:

- return a status code of Invalid Field in Command; and
- add an entry to the Error Information log page (refer to section 5.16.1.2).

...

Modify section 8 as shown below:

8 Extended Capabilities

...

8.3 Capacity Management

8.3.1 Overview

...

If an NVM Set is deleted, then the controller performs the following actions in sequence:

- 1) the NVM Set Identifier is removed from the NVM Set List;
- 2) if the Media Unit Status log page is supported, then the NVM Set Identifier field is cleared to 0h in all Media Unit Status Descriptors, if any, that indicated the deleted NVM Set;
- 3) for each namespace in the deleted NVM Set:
 - a. all commands targeting the namespace are handled as described for namespace deletion in section 8.11;
 - b. the namespace identifier is removed from the Allocated Namespace ID list;

Technical input submitted to the NVM Express® Workgroup is subject to the terms of the NVM Express® Participant's agreement. Copyright © 2008 to 2024 NVM Express, Inc.

- c. the namespace is deleted;
- d. the namespace identifier is added to the Changed Namespace List; and
- e. a Namespace Attribute Changed event is generated for hosts other than the host which issued the Capacity Management command;

and

- 4) the Unallocated Endurance Group Capacity indicated by the Endurance Group Information log page is increased by the amount of capacity formerly allocated to the NVM Set.

...

8.7 Directives

...

Support for a specific Directive Type is indicated using the Return Parameters operation of the Identify Directive. A specific Directive may be enabled or disabled using the Enable operation of the Identify Directive. Before using a specific Directive, the host should determine if that Directive is supported and should enable that Directive using the Identify Directive.

...

8.7.2 Identify (Directive Type 00h)

...

8.7.2.2 Directive Send

...

8.7.2.2.1 Enable Directive (Directive Operation 01h)

...

Figure 420: Enable Directive – Command Dword 12

Bits	Description
31:16	Reserved
15:08	Directive Type (DTYPE): This field specifies the Directive Type to enable or disable. If this field specifies the Identify Directive (i.e., 00h), then a status code of Invalid Field in Command shall be returned.
07:01	Reserved
00	Enable Directive (ENDIR): If set to '1' and the Directive Type is supported, then the Directive is enabled. If cleared to '0', then the Directive is disabled. If this bit is set to '1' for a Directive that is not supported, then a status code of Invalid Field in Command shall be returned.

Figure NEW: Enable Directive – Command Specific Status Values

Value	Description
7Fh	Stream Resource Allocation Failed: The controller failed to enable the Streams Directive for the specified namespace and Host Identifier due to insufficient resources.

...

8.12 Namespace Write Protection

...

Figure 429 defines the write protection states that may be supported for a namespace. All states persist across power cycles and Controller Level Resets (refer to section 3.7.2) except Write Protect Until Power Cycle state, which transitions to the No Write Protect state on the occurrence of a power cycle.

Technical input submitted to the NVM Express® Workgroup is subject to the terms of the NVM Express® Participant's agreement. Copyright © 2008 to 2024 NVM Express, Inc.

Figure 429: Namespace Write Protection State Definitions

State	Definition	Persistent Across	
		Power Cycles	Controller Level Resets
No Write Protect	The namespace is not write protected.	Yes	Yes
Write Protect	The namespace is write protected.	Yes	Yes
Write Protect Until Power Cycle	The namespace is write protected until the next power cycle.	No	Yes
Permanent Write Protect	The namespace is permanently write protected.	Yes	Yes

The Write Protect Until Power Cycle state should not be used in multi-domain NVM subsystems because clearing that state requires simultaneous power cycle of the namespace and all controllers to which that namespace is attached. The result of a command that attempts to use that state in a multi-domain NVM subsystem is specified in section 5.27.1.28.

<Note to Editor: Change above “should not” to “shall not” in next numbered (non-lettered) release>:

...

8.26 Virtualization Enhancements

...

Flexible Resources are controller resources that may be assigned to the primary controller or one of its secondary controllers. The Virtualization Management command is used to provision the Flexible Resources between a primary controller and one of its secondary controller(s). A primary controller’s allocation of Flexible Resources may be modified using the Virtualization Management command and the change takes effect after any Controller Level Reset other than a Controller Reset (~~i.e., CC.EN transitions from ‘1’ to ‘0’~~). A secondary controller only supports having Flexible Resources assigned or removed when in the Offline state.

...

Modify section 9 as shown below:

9 Error Reporting and Recovery

...

9.4 Internal Controller Error Handling

~~Errors such as a DRAM failure or power loss notification indicate that~~ If a controller level failure (e.g., a DRAM failure) has occurred during the processing of a command, then the controller should abort the command with a status code of ~~the completion queue entry should indicate an~~ Internal Error status code. ~~Host software shall ignore a~~ Any data transfer associated with the command may be incomplete or incorrect, and therefore any transferred data should not be used for any purpose other than error reporting or diagnosis. The host may choose to re-submit the command or indicate an error to the higher level software.

...

Description of NVM Express NVM Command Set Specification 1.0c changes

Modify section 2 as shown below:

2 NVM Command Set Model

...

2.1 Theory of operation

...

2.1.3 Fused Operation

...

2.1.3.1 Compare and Write

The Compare and Write fused operation compares the contents of the logical block(s) specified in the Compare command to the data stored at the indicated LBA range. If the compare operation is successful, then the LBA range is updated with the data provided in the Write command. If the ~~Compare~~ compare operation is not successful, then the controller shall abort the Write command operation-is-aborted with a status code of Command Aborted due to Failed Fused Command and the contents in the LBA range are not modified. If the ~~Write~~ write operation is not successful, the Compare command operation completion status is unaffected.

The LBA range, if used, shall be the same for the two commands. If the LBA ranges do not match, then the controller should abort the commands should-be-aborted with a status code of Invalid Field in Command;.

Note: To ensure the Compare and Write is an atomic operation in a multi-host environment, host software should ensure that the size of a Compare and Write fused operation is no larger than the ACWU/NACWU (refer to section 2.1.4) and that Atomic Boundaries are respected (refer to section 2.1.4.3). Controllers may abort a Compare and Write fused operation that is larger than ACWU/NACWU or that crosses an Atomic Boundary with ~~an error~~ a status code of Atomic Write Unit Exceeded.

...

Modify section 3 as shown below:

3 I/O Commands for the NVM Command Set

...

3.2 NVM Command Set Commands

...

3.2.2 Copy command

...

The data bytes in the LBAs specified by each Source Range Entry shall be copied to the destination LBA range in the same order those LBAs are listed in the Source Range entries (e.g., the LBAs specified by Source Range entry 0 are copied to the lowest numbered LBAs specified by the SDLBA field, the LBAs specified by Source Range entry 1 are copied to the next consecutively numbered LBAs after the LBAs

Technical input submitted to the NVM Express® Workgroup is subject to the terms of the NVM Express® Participant's agreement. Copyright © 2008 to 2024 NVM Express, Inc.

copied for Source Range entry 0~~-specified by the SDLBA field~~). The read operations and write operations used to perform the copy may operate sequentially or in parallel.

...

Modify section 5 as shown below:

5 Extended Capabilities

...

5.8 Command Set Specific Capability

5.8.1 Get LBA Status

...

Upon receiving an LBA Status Information Alert asynchronous event, the host should send one or more Get Log Page commands for Log Identifier 0Eh with the Retain Asynchronous Event bit set to '1' in order to read the entire LBA Status Information log page (refer to section 4.1.4.54).

...

Description of NVM Express Management Interface Specification

1.2c changes

Modify section 1 as shown below:

1 Introduction

...

1.8 Definitions

...

1.8.21 NVMe-MI Message

A type of MCTP Message that is defined by this specification in sections 3.1 and 4.1. [Refer to See](#) the MCTP IDs and Codes specification and the NVMe Management Messages over MCTP Binding Specification for more details on this type of MCTP Message (note that NVMe-MI Messages are referred to as NVM Express Management Messages over MCTP in these specifications).

...

Modify section 3 as shown below:

3 Message Transport

...

3.1 NVMe-MI Messages

...

In the out-of-band mechanism, an NVMe-MI Message consists of the payload of one or more MCTP packets. The maximum sized NVMe-MI Message is 4,224 bytes (i.e., 4 KiB + 128 bytes). Refer to the NVMe Management Messages over MCTP Binding Specification. NVMe-MI Messages with lengths greater than 4,224 bytes are considered invalid NVMe-MI Messages. [Refer to See](#) section 4.2 for details on how NVMe-MI Messages are used in the out-of-band mechanism.

In the in-band tunneling mechanism, NVMe-MI Messages are not split into MCTP packets and the maximum NVMe-MI message size is equal to the Maximum Data Transfer Size (refer to the NVM Express Base Specification). [Refer to See](#) section 4.3 for details on how NVMe-MI Messages are used in the in-band tunneling mechanism.

...

3.1.1 Message Fields

...

3.1.1.1 Message Integrity Check

...

[Refer to See](#) section 4.2.1.5 for special requirements on how to construct the Response Message when the Management Controller issues a Replay Control Primitive with a non-zero Response Replay Offset.

Technical input submitted to the NVM Express® Workgroup is subject to the terms of the NVM Express® Participant's agreement. Copyright © 2008 to 2024 NVM Express, Inc.

...

Modify section 5 as shown below:

5 Management Interface Command Set

...

5.3 Controller Health Status Poll

...

Figure 81: Controller Health Data Structure (CHDS)

Bytes	Description
...	
05:04	Composite Temperature (CTEMP): This field contains a value corresponding to a temperature in Kelvins that represents the current composite temperature of the Controller and Namespace(s) associated with that Controller. The value of this field corresponds to the value in the Controller's SMART / Health Information Log page .
06	Percentage Used (PDLU): This field contains a vendor specific estimate of the percentage of NVM Subsystem life used based on the actual usage and the manufacturer's prediction of NVM life. The value of this field corresponds to the value in the Controller's SMART / Health Information Log page .
07	Available Spare (SPARE): This field contains a normalized percentage (0% to 100%) of the remaining spare capacity available. The value of this field corresponds to the value in the Controller's SMART / Health Information Log page .
08	Critical Warning (CWARN): This field indicates critical warnings for the state of the Controller. The value of this field corresponds to the value in the Controller's SMART / Health Information Log page .
...	...
...	

...

5.6 NVM Subsystem Health Status Poll

...

Figure 91: NVM Subsystem Health Data Structure (NSHDS)

Bytes	Description
...	
1	Smart SMART Warnings (SW): This field contains the Critical Warning field (byte 0) of the NVMe SMART / Health Information log page . Each bit in this field is inverted from the NVMe Express Base Specification definition (i.e., the management interface shall indicate a '0' value while the corresponding bit is set to '1' in the log page). Refer to the NVMe Express Base Specification for bit definitions.
2	Composite Temperature (CTEMP): This field indicates the current temperature in degrees Celsius. If a temperature value is reported, it should be the same temperature as the Composite Temperature from the SMART / Health Information log page of hottest Controller in the NVM Subsystem. The reported temperature range is vendor specific and shall not exceed the range -60 °C to +127 °C. The 8-bit format of the data is shown below. This field should not report a temperature that is older than 1 s. If recent data is not available, the Responder should indicate a value of 80h for this field.
...	...

Figure 91: NVM Subsystem Health Data Structure (NSHDS)

Bytes	Description
3	Percentage Drive Life Used (PDLU): Contains a vendor specific estimate of the percentage of NVM Subsystem NVM life used based on the actual usage and the manufacturer's prediction of NVM life. If an NVM Subsystem has multiple Controllers, the highest value is returned. A value of 100 indicates that the estimated endurance of the NVM in the NVM Subsystem has been consumed but may not indicate an NVM Subsystem failure. The value is allowed to exceed 100. Percentages greater than 254 shall be represented as 255. This value should be updated once per power-on hour and equal the Percentage Used value in the NVMe SMART / Health Information Log P page.
...	

...

Modify section 6 as shown below:

6 NVM Express Admin Command Set

...

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
...				
Get Log Page ²	02h	M	O	NVM Express Base Specification
...				
Notes:				
<ol style="list-style-type: none"> 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 in-band. 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—The Management Endpoint shall ignore the Retain Asynchronous Event (RAE) bit in the Get Log Page command (refer to the NVM Express Base Specification). If the RAE bit is supported and the log page is used with Asynchronous Events, then the Get Log Page command shall be processed as if the RAE bit is set to '1'. 				

...

Modify section 8 as shown below:

8 Management Architecture

...

8.2 Vital Product Data

...

Technical input submitted to the NVM Express® Workgroup is subject to the terms of the NVM Express® Participant's agreement. Copyright © 2008 to 2024 NVM Express, Inc.

8.2.5 Topology MultiRecord Area

...

8.2.5.2 Upstream Connector Element Descriptor

...

Figure 160: Upstream Connector Element Descriptor

Bytes	Factory Default	Description
...		
03	Impl Spec	Form Factor: This field indicates the Form Factor of the NVMe Storage Device. Refer to See Figure 161 for a list of defined values.
...		

...

Appendix A Technical Note: NVM Express Basic Management Command

...

Figure 176: Subsystem Management Data Structure

Command Code	Offset (byte)	Description
0	...	
	02	SMART Warnings: This field shall contain the Critical Warning field (byte 0) of the NVMe SMART / Health Information log page. Each bit in this field shall be inverted from the NVMe definition (i.e., the management interface shall indicate a '0' value while the corresponding bit is '1' in the log page). Refer to the NVM Express Base Specification for bit definitions. If there are multiple Controllers in the NVM Subsystem, the Management Endpoint shall combine the Critical Warning field from every Controller such that a bit in this field is: <ul style="list-style-type: none">Cleared to '0' if any Controller in the NVM Subsystem indicates a critical warning for that corresponding bit.Set to '1' if all Controllers in the NVM Subsystem do not indicate a critical warning for the corresponding bit.
	03	Composite Temperature (CTemp): This field indicates the current temperature in degrees Celsius. If a temperature value is reported, it should be the same temperature as the Composite Temperature from the SMART / Health Information log page of hottest Controller in the NVM Subsystem. The reported temperature range is vendor specific, and shall not exceed the range -60 °C to +127°C. This field should not report a stale temperature, which means that it was sampled more than 5 s prior. If recent data is not available, the Management Endpoint should indicate a value of 80h for this field. The field values are shown below. ...
	04	Percentage Drive Life Used (PDLU): Contains a vendor specific estimate of the percentage of NVM Subsystem NVM life used based on the actual usage and the manufacturer's prediction of NVM life. If an NVM Subsystem has multiple Controllers the highest value is returned. A value of 100 indicates that the estimated endurance of the NVM in the NVM Subsystem has been consumed but may not indicate an NVM Subsystem failure. The value is allowed to exceed 100. Percentages greater than 254 shall be represented as 255. This value should be updated once per power-on hour and equal the Percentage Used value in the NVMe SMART / Health Information Log Page.

Technical input submitted to the NVMe Express® Workgroup is subject to the terms of the NVMe Express® Participant's agreement. Copyright © 2008 to 2024 NVMe Express, Inc.

Figure 176: Subsystem Management Data Structure

Command Code	Offset (byte)	Description
	...	
...		

...

Description of NVM Express Key Value Command Set Specification

1.0c changes

Modify section 3 as shown below:

3 I/O Commands for the Key Value Command Set

...

3.1 Submission Queue Entry and Completion Queue Entry

...

3.1.2 Key Value Command Set Specific Status Values

No command specific status values are defined in this specification.

This specification supports the Generic Command Specific status values defined in the NVM Express Base Specification. Generic Command Specific status values that are reported by commands defined in this specification are described ~~specific to the Key Value Command Set are defined~~ in Figure 4.

Figure 4: Status Code – Generic Command Specific Status Values, Key Value Command Set

Value	Description	Commands Affected
81h	Capacity Exceeded	Store
82h	Namespace Not Ready	Delete, Exist, Retrieve, Store
83h	Reservation Conflict	Delete, Store, Retrieve
84h	Format In Progress	Delete, Exist, List, Retrieve, Store
85h	Invalid Value Size	Store
86h	Invalid Key Size	List, Retrieve, Store
87h	KV Key Does Not Exist	Delete, Exist, Retrieve, Store
88h	Unrecovered Error	Retrieve
89h	Key Exists	Store

...

3.2 Key Value Command Set Commands

...

3.2.1 Delete command

...

3.2.1.1 Command Completion

Upon completion of the Delete command, the controller posts a completion queue entry (CQE) to the associated I/O Completion Queue. If the status code returned is 00h, then the KV key and its associated KV value have been deleted.

Delete command generic specific status values are defined in Figure 9.

< Note to Editor: Value 87h description title needs to be made bold, as shown below. >

Figure 9: Delete – Generic Command Specific Status Values

Value	Description
87h	KV Key Does Not Exist: The KV key does not exist.
0Bh	Invalid Namespace or Format: The namespace or the format of that namespace is invalid or the namespace is not associated with the Key Value Command Set.

3.2.2 List command

...

3.2.2.1 Command Completion

Upon completion of the List command, the controller shall post a completion queue entry to the associated I/O Completion Queue indicating the status for the command.

The command returns a list of KV keys that exist as described in 3.2.2.2.

List command generic specific status values are defined in Figure 14.

Figure 14: List – Generic Command Specific Status Values

Value	Description
86h	Invalid Key Size: The KV key size is not valid.
0Bh	Invalid Namespace or Format: The namespace or the format of that namespace is invalid.

...

3.2.3 Retrieve command

...

3.2.3.1 Command Completion

...

Retrieve command generic specific status values are defined in Figure 22.

Figure 22: Retrieve – Generic Command Specific Status Values

Value	Description
86h	Invalid Key Size: The KV key size is not valid.
0Bh	Invalid Namespace or Format: The namespace or the format of that namespace is invalid.
87h	KV Key Does Not Exist: The KV key does not exist.
88h	Unrecovered Error: There was an unrecovered error when reading from the medium.

3.2.4 Exist command

...

3.2.4.1 Command Completion

Upon completion of the Exist command, the controller posts a completion queue entry (CQE) to the associated I/O Completion Queue. If the status code returned is 00h, then the KV key exists.

The Exist command generic specific status values are defined in Figure 26.

Figure 26: Exist – Generic Command Specific Status Values

Value	Description
87h	KV Key Does Not Exist: The KV key does not exist.

3.2.5 Store command

...

3.2.5.1 Command Completion

Upon completion of the Store command, the controller shall post a completion queue entry to the associated I/O Completion Queue indicating the status for the command.

Store command generic specific errors are defined in Figure 32.

Figure 32: Store – Generic Command Specific Status Values

Value	Description
85h	Invalid Value Size: The value size is not valid.
86h	Invalid Key Size: The KV key size is not valid.
0Bh	Invalid Namespace or Format: The namespace or the format of that namespace is invalid.
81h	Capacity Exceeded: The capacity of the device was exceeded.
89h	Key Exists: Store Option bit 9 is set to '1' and the KV key exists.
87h	KV Key Does Not Exist: Store Option bit 8 is set to '1' and the KV key does not exist.

...