



LEGAL NOTICE:

© **Copyright 2008 to 2022 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 2022 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 2022 NVM Express, Inc.

NVM Express® Technical Errata

Errata ID	109
Revision Date	08/05/2022
Affected Spec Ver.	NVM Express® Base Specification Revision 2.0b NVM Express® NVM Command Set Specification Revision 1.0b NVM Express® Zoned Namespace Command Set Specification Revision 1.1b NVM Express® Key Value Command Set Specification Revision 1.0b NVM Express® Management Interface Specification 1.2b
Corrected Spec Ver.	

Errata Author(s)

Name	Company
Judy Brock, Mike Allison, Bill Martin	Samsung
Austin Bolen	DellEMC
Paul Suhler	Kioxia
Walt Hubis	Micron
Yoni Shternhell	Western Digital

Errata Overview

This ECN updates and clarifies various text within the NVM Express Base Specification Revision 2.0b, NVM Express NVM Command Set Specification Revision 1.0b, and NVM Express Key Value Command Set Specification Revision 1.0a.

Revision History

Revision Date	Change Description
10/31/2021	Initial creation
2/1/2022	Indicated that vendor specific LIDS and FIDs and option in the NVMe-MI.
2/2/2022	Updated capitalization of Feature.
2/3/2022	Added a backlog of outstanding ECN issues provided by Fred Knight.
2/7/2022	Cleanup of table rows not modified. Removed the word "Field" from field names. Added clarify when no protection information checking occurs.
2/8/2022	LBA Format 0 Extension definition. Fixed grammar on the timing text changes.
2/9/2022	Updated editorial items in section 5.3 of the NVM Command Set Specification.
2/14/2022	Updated comments from Yoni Shternhell.
2/16/2022	Updated navigation pane for section 4.1.5.1 in the NVM Command Set Specification. Inserted "capability" into "Namespace Management content" since either "command" or "capability" is with capitalized text.
2/17/2022	Editorial changes from review in the Technical WG. Updated the packing of LBA formats and KV Formats to use the field name as the starting point. Removed section 3.5.2 in the NVMe Base Specification as that is covered by ECN105.

Revision Date	Change Description
2/28/2022	Added the word "Parameter" to the LID Specific field and the Log Specific field.
3/2/2022	Striked red text. Converted references and cross-references to text.
3/3/2022	Changed the wording "until after" to "by the" to correct the timing.
3/8/2022	Corrected timing text associated with TP4084. Aligned acronyms of fields to be specified in parenthesis the word "field".
3/10/2022	Edits during review. Ready for 30 day member review.
4/8/2022	30 day member resolution
4/14/2022	Removed a double "is" in a sentence.
07/06/2022	Integrated
07/10/2022	Editorial changes per Mike Allison
08/02/2022	Editorial changes per Mike Allison
08/05/2022	Editorial fix of figure number and reference.

Description of Changes

NVM Express Base Specification 2.0b:

Backward Incompatible Changes:

- Clarified that the Release Context Action does not return log page data.

Editorial Changes:

- Update text to add the word "command" to the use of Get Log Page.
- Clarified that namespaces are created or exist as opposed to being allocated.
- Removed "Log Page" from the heading of the Predictable Latency Event Aggregate Log Page.
- Updated the consistent use of lower case use of "log page".
- Updated the consistent use of upper case use of "Persistent Event" when referencing the log page.
- Clarified that the reset value of the NSSRO bit for a firmware activation is in the domain.
- Correct capitalization of Feature.
- Clarified that the Supported Log Pages log page returned data is dependent on the UUID Index field value when the Get Log Page command supports the UUID Index field.
- Clarified that the Supported Feature Identifiers Supported and Effects log page returned data is dependent on the UUID Index field value when the Get Log Page command supports the UUID Index field.
- Aligned the usage of non-volatile storage medium when referring to the word medium where appropriate.
- Changed the name of the "Log Specific Field" to "Log Specific Parameter".
- Changed the name of the "LID Specific Field" to "LID Specific Parameter"
- Changed the name of the "Status Field" to "Status" so references are Status field.
- Clarified the timing on the usage of CRYPTO.CRWM and CRYPTO.CRIMT.
- Removed all Features defined by the NVM Command Set Specification from the feature list for Administrative controllers and Discovery controllers.
- Clarified timing text by change the use of "up to and including" to use "until after" type language.
- Fixed usage of acronyms to fields to be before the word "field".

- Removed the reference to the Command Submission and Execution field in the Feature Identifiers Supported and Effects as that field does not exist.
- Updated Feature definitions to capitalize the word “Feature” when use to reference the specific feature being defined.
- Capitalized “admin queue” for consistency.
- Move the word “field” after acronym of field name in parenthesis for consistency.

NVM Express NVM Command Set Specification 1.0b:

Editorial Changes:

- Updated the consistent use of lower case use of “log page”.
- Clarified that LBA formats are reported as opposed to being allocated.
- Clarified that namespaces are created or exist as opposed to being allocated.
- Aligned the usage of non-volatile storage medium when referring to the word medium where appropriate.
- Removed the word “Field” from actual field names: Protection Information Field, Protection Information Field Write, Protection Information Field Read.
- Clarified that no checking of the protection information is performed if the namespace is formatted with protection information, the PRACT bit is set to ‘1’, and the namespace is formatted with Metadata Size greater than protection information size.
- Clarified that the Namespace Management section is referring to the Namespace Management capability.
- Changed the reference of the NVM Express Base Specification to the correct figure in the document in the Namespace Management section.

NVM Express Zoned Namespace Command Set 1.0b:

Editorial Changes:

- Added the Invalid Protection Information status code of the Zone Append command.
- Removed the word “Field” from actual field names: Protection Information Field and Zone Receive Action Specific Field.
- Clarified the reference in the definition of the LBA Format 0 Extension.
- Clarified definition of Zone Descriptor Changed by adding a comma.

NVM Express Key Value Command Set 1.0b:

Editorial Changes:

- Clarified that KV formats are reported as opposed to being allocated.

NVM Express Management Interface Specification 1.0b:

Editorial Changes:

- Converted the Rocksoft™ parameters to a figure to allow it to split pages and retain the header rows defining the information (this aligns with the change in ECN108 in the NVM Command Set Specification).
- Clarified that the Max Payload Size setting is from the PCI Express Base Specification in the definition of the PCIe Maximum Payload Size field in the PCIe Port Specific Data.
- Updated the Log Page Support figure to clarify the requirements are for supporting the log pages.
- Editorial changes to correct pluralization.
- Added Vendor Specific log pages as optional.

- Added Vendor Specific Features as optional.
- Updates the reference to the NVM Express Base Specification change of the “Status Field” changing the name to “Status”.

Editor's Note:

BLACK text indicates unchanged text; **BLUE** text indicates newly inserted text, **RED** text indicates deleted text; **GREEN** text indicates editor notes.

Description of NVM Express Base Specification 2.0b changes

3 NVM Express Architecture

...

Description of NVM Express Base Specification 2.0b changes

Modify a portion of section 3.1.2.2.3 as shown below:

3.1 NVM Controller Architecture

...

3.1.2 Controller Types

...

3.1.2.2 Administrative Controller

...

3.1.2.2.3 Features Support

...

Figure 30: Administrative Controller – Feature Support

Feature Name	Feature Support Requirements ¹	Logged in Persistent Event Log ¹
...		
Error Recovery	P	P
...		
Write Atomicity Normal	P	P
...		

...

Modify a portion of section 3.1.2.3.4 as shown below:

3.1.2.3 Discovery Controller

...

3.1.2.3.4 Features Support

...

Figure 34: Discovery Controller – Feature Support

Feature Name	Feature Support Requirements ¹	Logged in Persistent Event Log ¹
...		
Error Recovery	P	P

Figure 34: Discovery Controller – Feature Support

Feature Name	Feature Support Requirements ¹	Logged in Persistent Event Log ¹
...		
Write Atomicity Normal	P	P
...		

...

Modify a portion of section 3.1.3.6 as shown below:

3.1.3 Controller Properties

...

3.1.3.6 Offset 1Ch: CSTS – Controller Status

Figure 47: Offset 1Ch: CSTS – Controller Status

Bits	Type	Reset ¹	Description
...			
04	RWC	HwInit	<p>NVM Subsystem Reset Occurred (NSSRO): The initial value of this bit is set to '1' if the last occurrence of an NVM Subsystem Reset (refer to section 3.7.1) occurred while power was applied to the domain. The initial value of this bit is cleared to '0' following an NVM Subsystem Reset due to application of power to the domain. This bit is only valid if the controller supports the NVM Subsystem Reset feature defined in section 3.7.1 as indicated by CAP.NSSRS set to '1'.</p> <p>The reset value of this bit is cleared to '0' if an NVM Subsystem Reset causes activation of a new firmware image in the domain.</p>
...			

...

Modify a portion of section 3.3.2.1.2 as shown below:

3.3 NVM Queue Models

...

3.3.2 Message-based Transport Queue Model

...

3.3.2.1 Capsules and Data Transfers

...

3.3.2.1.2 Response Capsules

...

Figure 82: Fabrics Response Capsule – Completion Queue Entry Format

Bytes	Description						
...							
15:14	Status (STS): Specifies status for the associated Fabrics command. <table> <tr> <th>Bits</th><th>Definition</th></tr> <tr> <td>15:01</td><td>Status Field as defined in section 3.3.3.2.1.</td></tr> <tr> <td>00</td><td>Reserved</td></tr> </table>	Bits	Definition	15:01	Status F ield as defined in section 3.3.3.2.1.	00	Reserved
Bits	Definition						
15:01	Status F ield as defined in section 3.3.3.2.1.						
00	Reserved						

...

Modify a portion of section 3.3.2.3 as shown below:

3.3.2.3 Queue Initialization and Queue State

...

The Authentication Requirements (AUTHREQ) field in the Connect response indicates if NVMe in-band authentication is required. If AUTHREQ is cleared to ~~0h~~zero, the created queue is ready for use after the Connect command completes successfully. If AUTHREQ is set to a non-zero value, the created queue is ready for use after NVMe in-band authentication has been performed successfully using the Authentication Send and Authentication Receive Fabrics commands.

...

Modify a portion of section 3.3.3.2 as shown below:

3.3.3 Queueing Data Structures

...

3.3.3.2 Common Completion Queue Entry

...

Figure 89: Common Completion Queue Entry Layout – Admin and All I/O Command Sets

	31	23	15	7	0
DW0	Command Specific				
DW1	Command Specific				
DW2	SQ Identifier		SQ Head Pointer		
DW3	Status-Field		P	Command Identifier	

...

Figure 91: Completion Queue Entry: DW 3

Bits	Description
31:17	Status-Field (SF): Indicates the status for the command that is being completed. Refer to section 3.3.3.2.1.
...	

...

Modify a portion of section 3.3.3.2.1 as shown below:

3.3.3.2.1 Status Field Definition

The Status **F**ield defines the status for the command indicated in the completion queue entry, defined in Figure 92.

A value of 0h for the Status **F**ield indicates a successful command completion, with no fatal or non-fatal error conditions. Unless otherwise noted, if a command fails to complete successfully for multiple reasons, then the particular status code returned is chosen by the vendor.

Figure 92: Completion Queue Entry: Status Field

Bits	Description
...	
29:28	<p>Command Retry Delay (CRD): If the DNR bit is cleared to '0' and the host has set the Advanced Command Retry Enable (ACRE) field to 1h in the Host Behavior Support feature (refer to section 5.27.1.18), then:</p> <ul style="list-style-type: none"> a) a 00b CRD value indicates a command retry delay time of zero (i.e., the host may retry the command immediately); and b) a non-zero CRD value selects a field in the Identify Controller data structure (refer to Figure 275) that indicates the command retry delay time: <ul style="list-style-type: none"> • a 01b CRD value selects the Command Retry Delay Time 1 (CRDT1) field; • a 10b CRD value selects the Command Retry Delay Time 2 (CRDT2) field; and • a 11b CRD value selects the Command Retry Delay Time 3 (CRDT3) field. <p>The host should not retry the command until at least the amount of time indicated by the selected field has elapsed. It is not an error for the host to retry the command prior to that time.</p> <p>If the DNR bit is set to '1' in the Status Field or the ACRE field is cleared to 0h in the Host Behavior Support feature, then this field is reserved.</p> <p>If the SCT and SC fields are cleared to 0h, then this field should be cleared to 00b.</p>
...	

...

Modify a portion of section 3.3.3.2.1.1 as shown below:

3.3.3.2.1.1 Generic Command Status Definition

...

Figure 94: Status Code – Generic Command Status Values

Value	Description	I/O Command Set Specific	I/O Command Set(s) ¹
...			

Figure 94: Status Code – Generic Command Status Values

Value	Description	I/O Command Set Specific	I/O Command Set(s) ¹
24h	<p>Admin Command Media Not Ready: The Admin command requires access to media and the media is not ready. The Do Not Retry bit indicates whether re-issuing the command at a later time may succeed. This status code shall only be returned:</p> <ul style="list-style-type: none"> a) for Admin commands; and b) if the controller is in Controller Ready Independent of Media mode (i.e., CC.CRIME is bit is set to '1'). <p>This status code shall not be returned with the Do Not Retry bit cleared to '0' after the amount of time indicated by the Controller Ready With Media Timeout field (CRTO.CRWMT) field after the controller is enabled (i.e., CC.EN transitions from '0' to '1').</p> <p>Refer to Figure 103 for the list of Admin commands permitted to return this status code.</p>	No	
...			

...

Modify a portion of section 3.5.3 as shown below:

3.5.3 Controller Ready Modes During Initialization

There are two possible controller ready modes:

- **Controller Ready With Media:** By the time the controller becomes ready (i.e., by the time that CSTS.RDY transitions from '0' to '1') after the controller is enabled (i.e., CC.EN transitions from '0' to '1'), then:
 - a) the controller shall be able to process all commands without error as described in section 3.5.4.1; and
 - b) all namespaces attached to the controller and all media required to process Admin commands shall be ready (i.e., commands are not permitted to be aborted with a status code of Namespace Not Ready with the Do Not Retry bit cleared to '0' or Admin Command Media Not Ready with the Do Not Retry bit cleared to '0').
- **Controller Ready Independent of Media:** After the controller is enabled, all namespaces attached to the controller and media required to process Admin commands may or may not become ready by the time the controller becomes ready. Any NVM command that specifies one or more namespaces attached to the controller is permitted to be aborted with a status code of Namespace Not Ready with the Do Not Retry bit cleared to '0' ~~for only up to and including until~~ CRTO.CRWMT amount of time after the controller is enabled.

Admin commands that require access to the media are permitted to be aborted with a status code of Admin Command Media Not Ready with the Do Not Retry bit cleared to '0' ~~for only up to and including until~~ CRTO.CRWMT amount of time after the controller is enabled. Refer to Figure 103 for a list of Admin commands that are permitted to be aborted with a status code of Admin Command Media Not Ready.

The controller shall be able to process without error as described in section 3.5.4.1-:

- a) all Admin commands not listed in Figure 103 by the time the controller is ready;
- b) all Admin commands listed in Figure 103 ~~no later than within, up to, and including~~ CRTO.CRWMT amount of time after the controller is enabled; and

- c) all NVM commands ~~no later than~~ ~~within, up to, and including~~ CRYPTO.CRWM amount of time after the controller is enabled.

...

Modify a portion of section 3.5.4 as shown below:

3.5.4 Controller Ready Timeouts During Initialization

...

3.5.4.1 Handling Errors During Initialization

If the CAP.CRMS field is non-zero and the controller has been enabled by transitioning CC.EN from '0' to '1' and the controller encounters a failure that prevents:

- a) at least one:

- command that does not access attached namespaces; or
- Admin command that does not require access to media (refer to Figure 103),

from being able to be processed without error by the amount of time indicated by the:

- Controller Ready Independent of Media Timeout (CRYPTO.CRIMT) field since the controller was enabled if the controller is in Controller Ready Independent of Media mode (i.e., the CC.CRIME bit is set to '1'); or
 - Controller Ready With Media Timeout (CRYPTO.CRWMT) field since the controller was enabled if the controller is in Controller Ready With Media mode (i.e., the CC.CRIME bit is cleared to '0');
- b) at least one namespace attached to the controller from becoming ready ~~within, up to, and including~~ by the amount of time indicated by the Controller Ready With Media Timeout (CRYPTO.CRWMT) field ~~amount of time since~~ after the controller was enabled; or
- c) media required by at least one Admin command from becoming ready ~~within, up to, and including~~ by the amount of time indicated by the Controller Ready With Media Timeout (CRYPTO.CRWMT) field ~~amount of time since~~ after the controller was enabled,

then:

- a) if the controller has not become ready, then the controller shall become ready (i.e., set CRYPTO.RDY to '1') ~~within up to and including~~ no later than CRYPTO.CRWMT amount of time after the controller was enabled; and
- b) if the Persistent Event Log page is supported, then the controller shall record an NVM Subsystem Hardware Error Event with the NVM Subsystem Hardware Error Event code set to a value of Controller Ready Timeout Exceeded in the Persistent Event Log page (refer to Figure 233).

Modify a portion of section 4.1.2 as shown below:

4 Data Structures

...

4.1 Data Layout

...

4.1.2 Scatter Gather List (SGL)

...

The value of the SDT field shall be less than or equal to the value of the Maximum SGL Data Block Descriptors ~~field~~ (MSDBD) field in the Identify Controller data structure (refer to the NVMe over Fabrics specification for the definition of the MSDBD field).

...

Modify a portion of section 5.2 as shown below:

5 Admin Command Set

...

5.2 Asynchronous Event Request command

...

5.2.1 Command Completion

...

Figure 146: Asynchronous Event Information – Notice

Value	Description
...	
03h	<p>Asymmetric Namespace Access Change: The Asymmetric Namespace Access information (refer to section 5.16.1.13) related to an ANA Group that contains namespaces attached to this controller has changed (e.g., an ANA state has changed, an ANAGRPID has changed). The current Asymmetric Namespace Access information for attached namespaces is indicated in the Asymmetric Namespace Access log page (refer to section 5.16.1.13). To clear this event, the host issues a Get Log Page command with the Retain Asynchronous Event bit cleared to '0' for the Asymmetric Namespace Access log.</p> <p>A controller shall not send this event if a Namespace Attribute Changed notice is sent for the same event, such as a change due to:</p> <ul style="list-style-type: none">a) the attachment of a namespace (refer to section 5.22);b) the deletion of a namespace (refer to section 5.23); orc) the detachment of a namespace (refer to section 5.22).
...	

...

Modify a portion of section 5.14 as shown below:

5.14 Format NVM command

...

Figure 188: Format NVM – Operation Scope

FNA Bit ¹	NSID	Format Operation
0	FFFFFFFFh ²	All namespaces attached to the controller. Other namespaces are not affected.
0	Any allocated value (refer to section 3.2.1.3)	Particular namespace specified. Other namespaces are not affected.
1 ³	Any allocated value (refer to section 3.2.1.3) or FFFFFFFFFh	All allocated namespaces that exist in the NVM subsystem. <add a period>

Figure 188: Format NVM – Operation Scope

FNA Bit ¹	NSID	Format Operation
...		

Modify a portion of section 5.16 as shown below:

5.16 Get Log Page command

...

Figure 197: Get Log Page – Command Dword 10

Bits	Description
...	
14:08	Log Specific Parameter Field (LSP): If not defined for the log specified by the Log Page Identifier field, this field is reserved.
...	

...

Modify a portion of section 5.16.1.1 as shown below:

5.16.1 Log Specific Information

...

5.16.1.1 Supported Log Pages (Log Identifier 00h)

An NVM subsystem may support several interfaces for submitting a Get Log Page command such as an Admin Submission Queue, PCIe VDM Management Endpoint, or SMBus/I2C Management Endpoint (refer to the NVM Express Management Interface Specification for details on Management Endpoints) and may have zero or more instances of each of those interfaces. The log pages supported on each instance of each interface may be different. This log page is used to describe the log pages that are supported on the interface to which the Get Log Page command was submitted and attributes specific to each log page. The log page is defined in Figure 203. The attributes of each log page are described in a LID Supported and Effects data structure defined in Figure 204.

If the UUID Selection Supported bit is set to '1' for the Get Log Page command in the Commands Supported and Effects log page (refer to section 5.16.1.6), then the log page data reflects the log pages that are supported based on the value of the UUID Index field (refer to section 8.25).

...

Figure 204: Get Log Page – LID Supported and Effects Data Structure

Bits	Description
31:16	LID Specific Parameter Field (LIDSP): This field is specific to the log page identifier as defined in Figure 205.
...	
...	

Figure 205: LID Supported and Effects Data Structure – LID Specific Parameter Field

Log Page Identifier	LID Specific Parameter Field						
...							
0Dh	<p>The LID Specific Parameter Field for log page identifier 0Dh (Persistent Event Log page as described in section 5.16.1.14) is defined as follows:</p> <table> <tr> <th>Bits</th><th>Description</th></tr> <tr> <td>15:1</td><td>Reserved</td></tr> <tr> <td>0</td><td> <p>Establish Context and Read 512 Bytes of Header Supported: If this bit is cleared to '0', then the controller does not support the Establish Context and Read 512 Bytes of Header action (refer to Figure 223). If this bit is set to '1', then the controller supports the Establish Context and Read 512 Bytes of Header action. If this bit is set to '1', then the Generation Number field in the Persistent Event Log page shall also be supported.</p> <p>Implementations compliant to later than NVM Express revision 1.4 shall set this bit to '1'.</p> </td></tr> </table>	Bits	Description	15:1	Reserved	0	<p>Establish Context and Read 512 Bytes of Header Supported: If this bit is cleared to '0', then the controller does not support the Establish Context and Read 512 Bytes of Header action (refer to Figure 223). If this bit is set to '1', then the controller supports the Establish Context and Read 512 Bytes of Header action. If this bit is set to '1', then the Generation Number field in the Persistent Event Log page shall also be supported.</p> <p>Implementations compliant to later than NVM Express revision 1.4 shall set this bit to '1'.</p>
Bits	Description						
15:1	Reserved						
0	<p>Establish Context and Read 512 Bytes of Header Supported: If this bit is cleared to '0', then the controller does not support the Establish Context and Read 512 Bytes of Header action (refer to Figure 223). If this bit is set to '1', then the controller supports the Establish Context and Read 512 Bytes of Header action. If this bit is set to '1', then the Generation Number field in the Persistent Event Log page shall also be supported.</p> <p>Implementations compliant to later than NVM Express revision 1.4 shall set this bit to '1'.</p>						
...							

...

Modify a portion of section 5.16.1.2 as shown below:

5.16.1.2 Error Information (Log Identifier 01h)

This log page is used to describe extended error information for a command that completed with error or report an error that is not specific to a particular command. Extended error information is provided when the More (M) bit is set to '1' in the Status Field for the completion queue entry associated with the command that completed with error or as part of an asynchronous event with an Error status type. This log page is global to the controller.

...

Figure 206: Get Log Page – Error Information Log Entry (Log Identifier 01h)

Bytes	Description						
...							
13:12	<table> <tr> <th>Bits</th><th>Description</th></tr> <tr> <td>15:1</td><td>Status Field: This field indicates the Status Field for the command that completed. If the error is not specific to a particular command, then this field reports the most applicable status value.</td></tr> <tr> <td>0</td><td>Phase Tag: This field may indicate the Phase Tag posted for the command.</td></tr> </table>	Bits	Description	15:1	Status Field: This field indicates the Status Field for the command that completed. If the error is not specific to a particular command, then this field reports the most applicable status value.	0	Phase Tag: This field may indicate the Phase Tag posted for the command.
Bits	Description						
15:1	Status Field: This field indicates the Status Field for the command that completed. If the error is not specific to a particular command, then this field reports the most applicable status value.						
0	Phase Tag: This field may indicate the Phase Tag posted for the command.						
...							

...

Modify a portion of section 5.16.1.8 as shown below:

5.16.1.8 Telemetry Host-Initiated (Log Identifier 07h)

This log consists of a header describing the log and zero or more Telemetry Data Blocks (refer to section 8.24). All Telemetry Data Blocks are 512 bytes in size. The controller shall initiate a capture of the controller's internal controller state to this log if the controller processes a Get Log Page command for this log with the Create Telemetry Host-Initiated Data bit set to '1' in the Log Specific Parameter field. 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 abort the command with a status code of Invalid Field in Command. This log page is global to the controller or global to the NVM subsystem.

Figure 214: Command Dword 10 – Log Specific Parameter Field

Bits	Description
...	

...

Modify a portion of section 5.16.1.12 as shown below:

5.16.1.12 Predictable Latency Event Aggregate Log Page (Log Identifier 0Bh)

...

If there is an enabled Predictable Latency Event pending for an NVM Set, then the Predictable Latency Event Aggregate log page includes an entry for that NVM Set. The log page is an ordered list by NVM Set Identifier. For example, if Predictable Latency Events are pending for NVM Set 27, 13, and 17, then the log page shall have entries in numerical order of 13, 17, and 27. A particular NVM Set is removed from this log page after the Get Log Page command is completed successfully with the Retain Asynchronous Event bit cleared to '0' for the Predictable Latency Per NVM Set log page for that NVM Set.

...

Modify a portion of section 5.16.1.13 as shown below:

5.16.1.13 Asymmetric Namespace Access (Log Identifier 0Ch)

...

Figure 220: Command Dword 10 – Log Specific Parameter Field

Bits	Description
...	

...

Modify a portion of section 5.16.1.14 as shown below:

5.16.1.14 Persistent Event Log (Log Identifier 0Dh)

The Persistent Event Log page contains information about significant events not specific to a particular command. The information in this log page shall be retained across power cycles and resets. NVM subsystems should be designed for minimal loss of event information upon power failure. This log consists of a header describing the log and zero or more Persistent Events (refer to section 5.16.1.14.1).

...

The Action field in the Log Specific Parameter Ffield (refer to Figure 223) specifies if-whether:

- a) A persistent event log reporting context is created at the start of processing this Get Log Page command and log page data, if any, is read from the log page data associated with that log reporting context;
- b) Log page data is read from the log page data associated with a preexisting log reporting context;
- or
- c) The persistent event log reporting context, if any, is released.

...

The controller should retain the persistent event log reporting context:

- a) Until the controller processes:
 - a) **aA** Get Log Page command requesting the Persistent Event **L**og page with the Action field set to 02h (i.e., Release Context);
 - b) **aA**n NVM Subsystem Reset; or
 - c) **aA** Controller Level Reset;
- or
- b) For a vendor specific time long enough to allow retrieval of the **P**ersistent **E**vent log page data.

...

Figure 223: Command Dword 10 – Log Specific **Parameter Field**

Bits	Description												
...													
09:08	<p>Action: This field specifies the action the controller shall take during processing this Get Log Page command.</p> <table> <tr> <th>Value</th><th>Definition</th></tr> <tr> <td>00b</td><td> <p>Read Log Data: Return Persistent Event log page data starting at the address indicated by the LPOU field and the LPOL field in the Get Log Page command. If the controller does not have a persistent event log reporting context, then the controller shall abort the command with a status code of Command Sequence Error.</p> </td></tr> <tr> <td>01b</td><td> <p>Establish Context and Read Log Data: The controller shall:</p> <ol style="list-style-type: none"> a) determine the length of the Persistent Event log page data; b) determine the set of events to report in the Persistent Event log page data; and • establish a persistent event log reporting context to store information describing the Persistent Event log page data to be reported and track Persistent Event log page data accesses. <p>After establishing a persistent event log reporting context, the controller shall return Persistent Event log page data starting at the address indicated by the LPOU field and the LPOL field in the Get Log Page command.</p> <p>If a persistent event log reporting context already exists, then the controller shall abort the command with a status code of Command Sequence Error.</p> </td></tr> <tr> <td>10b</td><td> <p>Release Context: The controller shall:</p> <ul style="list-style-type: none"> • release the persistent event log reporting context, if any. It is not an error if the controller does not have a persistent event log reporting context; and • not return any Persistent Event log page data (i.e., the controller ignores the NUMDU field, NUMDL field, LPOL field, and the LPOU field). </td></tr> <tr> <td>11b</td><td> <p>Establish Context and Read 512 Bytes of Header: The controller shall:</p> <ol style="list-style-type: none"> a) determine the length of the Persistent Event log page data; b) determine the set of events to report in the Persistent Event log page data; and c) if a reporting context does not already exist, then establish a persistent event log reporting context to store information describing the Persistent Event log page data to be reported and track Persistent Event log page data accesses. </td></tr> <tr> <td>....</td><td></td></tr> </table>	Value	Definition	00b	<p>Read Log Data: Return Persistent Event log page data starting at the address indicated by the LPOU field and the LPOL field in the Get Log Page command. If the controller does not have a persistent event log reporting context, then the controller shall abort the command with a status code of Command Sequence Error.</p>	01b	<p>Establish Context and Read Log Data: The controller shall:</p> <ol style="list-style-type: none"> a) determine the length of the Persistent Event log page data; b) determine the set of events to report in the Persistent Event log page data; and • establish a persistent event log reporting context to store information describing the Persistent Event log page data to be reported and track Persistent Event log page data accesses. <p>After establishing a persistent event log reporting context, the controller shall return Persistent Event log page data starting at the address indicated by the LPOU field and the LPOL field in the Get Log Page command.</p> <p>If a persistent event log reporting context already exists, then the controller shall abort the command with a status code of Command Sequence Error.</p>	10b	<p>Release Context: The controller shall:</p> <ul style="list-style-type: none"> • release the persistent event log reporting context, if any. It is not an error if the controller does not have a persistent event log reporting context; and • not return any Persistent Event log page data (i.e., the controller ignores the NUMDU field, NUMDL field, LPOL field, and the LPOU field). 	11b	<p>Establish Context and Read 512 Bytes of Header: The controller shall:</p> <ol style="list-style-type: none"> a) determine the length of the Persistent Event log page data; b) determine the set of events to report in the Persistent Event log page data; and c) if a reporting context does not already exist, then establish a persistent event log reporting context to store information describing the Persistent Event log page data to be reported and track Persistent Event log page data accesses. 	
Value	Definition												
00b	<p>Read Log Data: Return Persistent Event log page data starting at the address indicated by the LPOU field and the LPOL field in the Get Log Page command. If the controller does not have a persistent event log reporting context, then the controller shall abort the command with a status code of Command Sequence Error.</p>												
01b	<p>Establish Context and Read Log Data: The controller shall:</p> <ol style="list-style-type: none"> a) determine the length of the Persistent Event log page data; b) determine the set of events to report in the Persistent Event log page data; and • establish a persistent event log reporting context to store information describing the Persistent Event log page data to be reported and track Persistent Event log page data accesses. <p>After establishing a persistent event log reporting context, the controller shall return Persistent Event log page data starting at the address indicated by the LPOU field and the LPOL field in the Get Log Page command.</p> <p>If a persistent event log reporting context already exists, then the controller shall abort the command with a status code of Command Sequence Error.</p>												
10b	<p>Release Context: The controller shall:</p> <ul style="list-style-type: none"> • release the persistent event log reporting context, if any. It is not an error if the controller does not have a persistent event log reporting context; and • not return any Persistent Event log page data (i.e., the controller ignores the NUMDU field, NUMDL field, LPOL field, and the LPOU field). 												
11b	<p>Establish Context and Read 512 Bytes of Header: The controller shall:</p> <ol style="list-style-type: none"> a) determine the length of the Persistent Event log page data; b) determine the set of events to report in the Persistent Event log page data; and c) if a reporting context does not already exist, then establish a persistent event log reporting context to store information describing the Persistent Event log page data to be reported and track Persistent Event log page data accesses. 												
....													

...

Modify a portion of section 5.16.1.14.5 as shown below:

5.16.1.14.1 Persistent Event Log Events

...

5.16.1.14.1.5 NVM Subsystem Hardware Error Event (Event Type 05h)

...

Figure 233: NVM Subsystem Hardware Error Event Codes

Code	Description
...	
0Bh	Controller Ready Timeout Exceeded: Indicates that: a) the controller was not ready to process at least one command without error as described in section 3.5.4.1 within up to and including the amount of time indicated by: <ul style="list-style-type: none">the Controller Ready With Media Timeout field (CRTO.CRWMT) field in Controller Ready With Media mode (CC.CRIME is cleared to '0'); orthe Controller Ready Independent of Media Timeout field (CRTO.CRIMT) field in Controller Ready Independent of Media mode (CC.CRIME is set to '1'); since the controller was enabled; or b) at least one namespace attached to the controller or media required to process at least one Admin command was not ready within up to and including the amount of time indicated by the Controller Ready With Media Timeout field (CRTO.CRWMT) field since after the controller was enabled by transitioning CC.EN from '0' to '1'. Refer to Figure 235 for the format of the Additional Hardware Error Information field.
...	

...

5.16.1.14.1.8 Format NVM Completion Event (Event Type 08h)

A Format NVM Completion event shall be recorded in the Persistent Event Log at the completion of a Format NVM command that resulted in modification of the contents of the NVM.

The Format NVM Completion event shall set the Persistent Event Log Event Format Header:

- Event Type field to 08h; and
- Event Type Revision field to 01h.

Figure 238: Format NVM Completion Event Data Format (Event Type 08h)

Bytes	Description
03:00	Namespace Identifier: Contains the namespace identifier specified in the Format NVM command.
04	Smallest Format Progress Indicator: For a Format NVM command that formats a single namespace, this field contains the lowest numerical value that was available for reporting in the FPI field of the Identify Namespace data structure (i.e., if the format did not complete successfully and the FPI field is supported, then this field contains the percentage of the namespace that remained to be formatted at the time the Format NVM command completed, refer to Figure 280) during the format operation. For a Format NVM command that formats all namespaces this field shall be cleared to 0h.

Figure 238: Format NVM Completion Event Data Format (Event Type 08h)

Bytes	Description
05	Format NVM Status: Contains the status of the format operation.
07:06	Completion Information: Contains a vendor specific value that may provide more information about the completion of the format operation (e.g., if the format operation did not complete successfully, then this field may contain a vendor specific code that indicates a vendor specific reason).
11:08	Status: Contains the value that was reported in the status code field for the completion queue entry, if any, for the Format NVM command associated with this event. If no completion queue entry was reported, then this field shall be cleared to 0h.

Figure 235: Additional Hardware Error Information for Controller Ready Timeout Exceeded errors

Bytes	Value
0	Controller State: Indicates the state of the controller at the time the Controller Ready Timeout Exceeded error occurred.
3:1	Reserved

...

Modify a portion of section 5.16.1.14.1.8 as shown below:

5.16.1.14.1.8 Format NVM Completion Event (Event Type 08h)

...

Figure 238: Format NVM Completion Event Data Format (Event Type 08h)

Bytes	Description
...	
11:08	Status Field: Contains the value that was reported in the status code field for the completion queue entry, if any, for the Format NVM command associated with this event. If no completion queue entry was reported, then this field shall be cleared to 0h.

...

Modify a portion of section 5.16.1.15 as shown below:

5.16.1.15 Endurance Group Event Aggregate (Log Identifier 0Fh)

...

If there is an enabled Endurance Group Event pending for an Endurance Group, then the Endurance Group Event Aggregate log page includes an entry for that Endurance Group. The log page is an ordered list by Endurance Group Identifier. For example, if Endurance Group Events are pending for Endurance Group 2, 1, and 7, then the log page shall have entries in numerical order of 1, 2, and 7. A particular Endurance Group entry is removed from this log page after the Get Log Page [command](#) is completed successfully with the Retain Asynchronous Event bit cleared to '0' for the Endurance Group Information log page for that Endurance Group.

...

Modify a portion of section 5.16.1.18 as shown below:

5.16.1.18 Feature Identifiers Supported and Effects (Log Identifier 12h)

An NVM subsystem may support several interfaces for submitting a Get Log Page command such as an Admin Submission Queue, PCIe VDM Management Endpoint, or SMBus/I2C Management Endpoint (refer the NVM Express Management Interface Specification for details on Management Endpoints) and may have zero or more instances of each of those interfaces. The feature identifiers (FIDs) supported on each instance of each interface may be different. This log page describes the FIDs that are supported on the interface to which the Get Log Page command was submitted and the effects of those features on the state of the NVM subsystem. The log page is defined in Figure 255. Each Feature Identifier's effects are described in a FID Supported and Effects data structure defined in Figure 256.

If the UUID Selection Supported bit is set to '1' for the Get Log Page command in the Commands Supported and Effects log page (refer to section [5.16.1.6](#)), then the log page data reflects the FIDs that are supported based on the value of the UUID Index field (refer to section [8.25](#)).

...

The FID Supported and Effects data structure describes the effect of a Set Features command for the FID, including any optional features of the FID.

<<The Command Submission and Execution field is not supported by this Feature>>

~~If multiple hosts are connected to the NVM subsystem, then those hosts should coordinate their commands to meet the Command Submission and Execution requirements (refer to Figure 256). The details of this coordination are outside the scope of this specification.~~

...

Modify a portion of section 5.16.1.20 as shown below:

5.16.1.20 Command and Feature Lockdown (Log Identifier 14h)

This log page is used to indicate which commands and Set Features Feature Identifiers are supported to be prohibited from execution using the Command and Feature Lockdown capability (refer to section 8.4) and which commands are currently prohibited if received on an NVM Express controller Admin Submission Queue or received out-of-band on a Management Endpoint (refer to the NVM Express Management Interface Specification). This log page uses the Log Specific **Parameter Field** field (refer to Figure 259) and may use the UUID Index field in the Get Log Page command to specify the scope and content of the list returned in the Command and Feature Identifier List field of this log page. The UUID Index field may be used if the Scope field is set to 2h, allowing returning of vendor specific Set Features Feature Identifier lockdown information.

Figure 259: Command Dword 10 – Log Specific **Parameter Field**

Bits	Description
...	

...

Figure 260: Get Log Page – Command and Feature Lockdown Log

Bytes	Description																																								
0	<table><tr><th>Bits</th><th>Description</th></tr><tr><td>7:6</td><td>Reserved</td></tr><tr><td rowspan="5">5:4</td><td>Contents Selected (CS): This field in combination with the Scope Selected field indicates the contents of the Command and Feature Identifier List field in the log page. The Content Selected field is specified by the contents of the Contents field in the Log Specific Parameter Field field of the Get Log Page command.</td></tr><tr><td><table><tr><th>Value</th><th>Description</th></tr><tr><td>00b</td><td>List contains command opcodes or Set Features Feature Identifiers based on the Scope Selected field that are supported to be prohibited</td></tr><tr><td>01b</td><td>List contains command opcodes or Set Features Feature Identifiers based on the Scope Selected field that are currently prohibited if received on an NVM Express controller submission queue</td></tr><tr><td>10b</td><td>List contains command opcodes or Set Features Feature Identifiers based on 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></td></tr><tr><td rowspan="6">3:0</td><td>Scope Selected (SS): This field in combination with the Contents Selected field indicates what the Command and Feature Identifier List field contains in the log page. The Scope Selected field is specified by the contents of the Scope field in the Log Specific Parameter field of the Get Log Page command.</td></tr><tr><td><table><tr><th>Value</th><th>Description</th></tr><tr><td>0h</td><td>List contains Admin Command Set opcodes</td></tr><tr><td>1h</td><td>Reserved</td></tr><tr><td>2h</td><td>List contains Set Features Feature Identifiers</td></tr><tr><td>3h</td><td>List contains Management Interface Command Set opcodes</td></tr><tr><td>4h</td><td>List contains PCIe Command Set opcodes</td></tr><tr><td>5h to Fh</td><td>Reserved</td></tr></table></td></tr><tr><td colspan="2">...</td></tr><tr><td colspan="2">...</td></tr><tr><td colspan="2">...</td></tr></table>	Bits	Description	7:6	Reserved	5:4	Contents Selected (CS): This field in combination with the Scope Selected field indicates the contents of the Command and Feature Identifier List field in the log page. The Content Selected field is specified by the contents of the Contents field in the Log Specific Parameter Field field of the Get Log Page command.	<table><tr><th>Value</th><th>Description</th></tr><tr><td>00b</td><td>List contains command opcodes or Set Features Feature Identifiers based on the Scope Selected field that are supported to be prohibited</td></tr><tr><td>01b</td><td>List contains command opcodes or Set Features Feature Identifiers based on the Scope Selected field that are currently prohibited if received on an NVM Express controller submission queue</td></tr><tr><td>10b</td><td>List contains command opcodes or Set Features Feature Identifiers based on 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	Description	00b	List contains command opcodes or Set Features Feature Identifiers based on the Scope Selected field that are supported to be prohibited	01b	List contains command opcodes or Set Features Feature Identifiers based on the Scope Selected field that are currently prohibited if received on an NVM Express controller submission queue	10b	List contains command opcodes or Set Features Feature Identifiers based on the Scope field that are currently prohibited if received out-of-band on a Management Endpoint	11b	Reserved	3:0	Scope Selected (SS): This field in combination with the Contents Selected field indicates what the Command and Feature Identifier List field contains in the log page. The Scope Selected field is specified by the contents of the Scope field in the Log Specific Parameter field of the Get Log Page command.	<table><tr><th>Value</th><th>Description</th></tr><tr><td>0h</td><td>List contains Admin Command Set opcodes</td></tr><tr><td>1h</td><td>Reserved</td></tr><tr><td>2h</td><td>List contains Set Features Feature Identifiers</td></tr><tr><td>3h</td><td>List contains Management Interface Command Set opcodes</td></tr><tr><td>4h</td><td>List contains PCIe Command Set opcodes</td></tr><tr><td>5h to Fh</td><td>Reserved</td></tr></table>	Value	Description	0h	List contains Admin Command Set opcodes	1h	Reserved	2h	List contains Set Features Feature Identifiers	3h	List contains Management Interface Command Set opcodes	4h	List contains PCIe Command Set opcodes	5h to Fh	Reserved	
	Bits	Description																																							
	7:6	Reserved																																							
	5:4	Contents Selected (CS): This field in combination with the Scope Selected field indicates the contents of the Command and Feature Identifier List field in the log page. The Content Selected field is specified by the contents of the Contents field in the Log Specific Parameter Field field of the Get Log Page command.																																							
		<table><tr><th>Value</th><th>Description</th></tr><tr><td>00b</td><td>List contains command opcodes or Set Features Feature Identifiers based on the Scope Selected field that are supported to be prohibited</td></tr><tr><td>01b</td><td>List contains command opcodes or Set Features Feature Identifiers based on the Scope Selected field that are currently prohibited if received on an NVM Express controller submission queue</td></tr><tr><td>10b</td><td>List contains command opcodes or Set Features Feature Identifiers based on 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	Description	00b		List contains command opcodes or Set Features Feature Identifiers based on the Scope Selected field that are supported to be prohibited	01b	List contains command opcodes or Set Features Feature Identifiers based on the Scope Selected field that are currently prohibited if received on an NVM Express controller submission queue	10b	List contains command opcodes or Set Features Feature Identifiers based on the Scope field that are currently prohibited if received out-of-band on a Management Endpoint	11b	Reserved																												
		Value	Description																																						
		00b	List contains command opcodes or Set Features Feature Identifiers based on the Scope Selected field that are supported to be prohibited																																						
		01b	List contains command opcodes or Set Features Feature Identifiers based on the Scope Selected field that are currently prohibited if received on an NVM Express controller submission queue																																						
	10b	List contains command opcodes or Set Features Feature Identifiers based on the Scope field that are currently prohibited if received out-of-band on a Management Endpoint																																							
	11b	Reserved																																							
3:0	Scope Selected (SS): This field in combination with the Contents Selected field indicates what the Command and Feature Identifier List field contains in the log page. The Scope Selected field is specified by the contents of the Scope field in the Log Specific Parameter field of the Get Log Page command.																																								
	<table><tr><th>Value</th><th>Description</th></tr><tr><td>0h</td><td>List contains Admin Command Set opcodes</td></tr><tr><td>1h</td><td>Reserved</td></tr><tr><td>2h</td><td>List contains Set Features Feature Identifiers</td></tr><tr><td>3h</td><td>List contains Management Interface Command Set opcodes</td></tr><tr><td>4h</td><td>List contains PCIe Command Set opcodes</td></tr><tr><td>5h to Fh</td><td>Reserved</td></tr></table>	Value	Description	0h	List contains Admin Command Set opcodes	1h	Reserved	2h	List contains Set Features Feature Identifiers	3h	List contains Management Interface Command Set opcodes	4h	List contains PCIe Command Set opcodes	5h to Fh	Reserved																										
	Value	Description																																							
	0h	List contains Admin Command Set opcodes																																							
	1h	Reserved																																							
	2h	List contains Set Features Feature Identifiers																																							
3h	List contains Management Interface Command Set opcodes																																								
4h	List contains PCIe Command Set opcodes																																								
5h to Fh	Reserved																																								
...																																									
...																																									
...																																									

...

Modify a portion of section 5.16.1.21 as shown below:

5.16.1.21 Boot Partition (Log Identifier 15h)

The Boot Partition Log page provides read only access to the Boot Partition (refer to section 8.2) accessible by this controller through the BPRSEL register (refer to section 3.1.3.14).

This log consists of a header describing the Boot Partition and Boot Partition data as defined by Figure 262. The Boot Partition Identifier bit in the Log Specific Parameter Field determines the Boot Partition.

...

Figure 261: Command Dword 10 – Log Specific Parameter Field

Bits	Description
...	

...

Modify a portion of section 5.27 as shown below:

5.27 Set Features command

...

Figure 314: Set Features – Command Dword 10

Bits	Description
31	Save (SV): This bit specifies that the controller shall save the attribute so that the attribute persists through all power states and resets. The controller indicates in bit 4 of the Optional NVM Command Support field of the Identify Controller data structure in Figure 275 whether this bit is supported. If the Feature Identifier specified in the Set Features command is not saveable by the controller and the controller receives a Set Features command with the Save bit set to '1'one, then the command shall be aborted with a status code of Feature Identifier Not Saveable.
...	

...

Modify a portion of section 5.27.1.3 as shown below:

5.27.1 Feature Specific Information

...

5.27.1.3 Temperature Threshold (Feature Identifier 04h)

A controller may report up to nine temperature values in the SMART / Health Information log page (i.e., the Composite Temperature and Temperature Sensor 1 through Temperature Sensor 8; refer to Figure 207). Associated with each implemented temperature sensor is an over temperature threshold and an under temperature threshold. When a temperature is greater than or equal to its corresponding over temperature threshold or less than or equal to its corresponding under temperature threshold, then bit 1one of the Critical Warning field in the SMART / Health Information Log page (refer to section 5.16.1.3) is set to '1'one. This may trigger an asynchronous event.

...

If a Get Features command is submitted for this Feature, the temperature threshold selected by Command Dword 11 is returned in Dword 0 of the completion queue entry for that command.

...

Modify a portion of section 5.27.1.11 as shown below:

5.27.1.11 Timestamp (Feature Identifier 0Eh), (Optional)

The Feature <<From TP4074a>> enables the host to set a timestamp value in the controller. A controller indicates support for the Timestamp feature through the Optional NVM Command Support (ONCS) field in the Identify Controller data structure. The Timestamp field value (refer to Figure 339) in a Set Features command sets a timestamp value in the controller. After the current value for this Feature is set, the controller updates that value as time passes. A Get Features command that requests the current value reports the timestamp value in the controller at the time the Get Features command is processed (e.g., the value set with a Set Features command for the current value plus the elapsed time since being set).

...

Modify a portion of section 5.27.1.13 as shown below:

5.27.1.13 Host Controlled Thermal Management (Feature Identifier 10h), (Optional)

...

Figure 342: HCTM – Command Dword 11

Bits	Description
31:16	<p>Thermal Management Temperature 1 (TMT1): This field specifies the temperature, in Kelvins, when the controller begins to transition to lower power active power states or performs vendor specific thermal management actions while minimizing the impact on performance (e.g., light throttling) in order to attempt to reduce the Composite Temperature.</p> <p>A value cleared to 0h, specifies that this part of the Feature shall be disabled.</p> <p>The range of values that are supported by the controller are indicated in the Minimum Thermal Management Temperature field and Maximum Thermal Management Temperature field in the Identify Controller data structure in Figure 275.</p> <p>If the host attempts to set this field to a value less than the value contained in the Minimum Thermal Management Temperature field or greater than the value contained in the Maximum Thermal Management Temperature field in the Identify Controller data structure in Figure 275, then the command shall abort with a status code of Invalid Field in Command.</p> <p>If the host attempts to set this field to a value greater than or equal to the value contained in the Thermal Management Temperature 2 field, if non-zero, then the command shall abort with a status code of Invalid Field in Command.</p>
15:00	<p>Thermal Management Temperature 2 (TMT2): This field specifies the temperature, in Kelvins, when the controller begins to transition to lower power active power states or perform vendor specific thermal management actions regardless of the impact on performance (e.g., heavy throttling) in order to attempt to reduce the Composite Temperature.</p> <p>A value cleared to 0h, specifies that this part of the Feature shall be disabled.</p> <p>The range of values that are supported by the controller are indicated in the Minimum Thermal Management Temperature field and Maximum Thermal Management Temperature field in the Identify Controller data structure in Figure 275.</p> <p>If the host attempts to set this field to a value less than the value contained in the Minimum Thermal Management Temperature field or greater than the value contained in the Maximum Thermal Management Temperature field in the Identify Controller data structure in Figure 275, then the command shall abort with a status code of Invalid Field in Command.</p> <p>If the host attempts to set this field to a non-zero value less than or equal to the value contained in the Thermal Management Temperature 1 field, then the command shall abort with a status code of Invalid Field in Command.</p>

...

Modify a portion of section 5.27.1.16 as shown below:

5.27.1.16 Predictable Latency Mode Config (Feature Identifier 13h)

...

Figure 348: Predictable Latency Mode – Deterministic Threshold Configuration Data Structure

Bytes	Description														
01:00	Enable Event: This field specifies whether an entry shall be added to the Predictable Latency Event Aggregate L og P age for the associated event. If a bit is set to '1', then an entry shall be added if the specified event occurs. If a bit is cleared to '0', then an entry shall not be added if the specified event occurs.														
	<table><tr><th>Bits</th><th>Description</th></tr><tr><td>00</td><td>DTWIN Reads Warning</td></tr><tr><td>01</td><td>DTWIN Writes Warning</td></tr><tr><td>02</td><td>DTWIN Time Warning</td></tr><tr><td>03 to 13</td><td>Reserved</td></tr><tr><td>14</td><td>Autonomous transition from DTWIN to NDWIN due to typical or maximum value exceeded.</td></tr><tr><td>15</td><td>Autonomous transition from DTWIN to NDWIN due to Deterministic Excursion.</td></tr></table>	Bits	Description	00	DTWIN Reads Warning	01	DTWIN Writes Warning	02	DTWIN Time Warning	03 to 13	Reserved	14	Autonomous transition from DTWIN to NDWIN due to typical or maximum value exceeded.	15	Autonomous transition from DTWIN to NDWIN due to Deterministic Excursion.
	Bits	Description													
	00	DTWIN Reads Warning													
	01	DTWIN Writes Warning													
	02	DTWIN Time Warning													
	03 to 13	Reserved													
14	Autonomous transition from DTWIN to NDWIN due to typical or maximum value exceeded.														
15	Autonomous transition from DTWIN to NDWIN due to Deterministic Excursion.														
31:02	Reserved														
39:32	DTWIN Reads Threshold: If the value of DTWIN Reads Estimate falls below this value and the DTWIN Reads Warning is enabled, then the 'DTWIN Reads Warning' event is set in the Predictable Latency Per NVM Set L og P age for the affected NVM Set.														
47:40	DTWIN Writes Threshold: If the value of DTWIN Writes Estimate falls below this value and the DTWIN Writes Warning is enabled, then the 'DTWIN Writes Warning' event is set in the Predictable Latency Per NVM Set L og P age for the affected NVM Set.														
55:48	DTWIN Time Threshold: If the value of DTWIN Time Estimate falls below this value and the DTWIN Time Warning is enabled, then the 'DTWIN Time Warning' event is set in the Predictable Latency Per NVM Set L og P age for the affected NVM Set.														
511:56	Reserved														

...

Modify a portion of section 5.27.1.22 as shown below:

5.27.1.22 Spinup Control (Feature Identifier 1Ah)

...

If the NVM subsystem does not contain any Endurance Groups that store data on rotational media, then the controller shall abort the Set Features command and the Get Features command for this Feature with status code of Invalid Field In Command.

...

Modify a portion of section 6.3 as shown below:

6 Fabrics Command Set

...

6.3 Connect Command and Response

...

The Connect response provides status for the Connect command. If a connection is established, then the Controller ID allocated to the host is returned. The Connect response is defined in Figure 382.

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.

...

Modify a portion of section 8.2 as shown below:

8 Extended Capabilities

...

8.2 Boot Partitions

Boot Partitions provide an optional area of NVM storage that may be read by the host without the host initializing queues or enabling the controller. The simplified interface to access Boot Partitions may be used for platform initialization code (e.g., a bootloader that is executed from host ROM) to boot to a pre-OS environment (e.g., UEFI) instead of storing the image on another [non-volatile](#) storage medium (e.g., SPI flash). Refer to section 8.2.1 for the procedure to read the contents of a Boot Partition.

...

Modify a portion of section 8.10.2 as shown below:

8.10 Host Operation with Asymmetric Namespace Access Reporting (Informative)

...

8.10.2 Host ANA Inaccessible Operation

If the ANA [Log](#) page reports an ANA state of ANA Inaccessible State for an ANA Group or a command returns a status code of Asymmetric Access Inaccessible, then the host should:

- not use that controller to send commands to any namespace in that ANA Group; and
- select a different controller for sending commands to all namespaces in that ANA Group.

...

Modify a portion of section 8.11 as shown below:

8.11 Namespace Management

...

Reporting of capacity information for the NVM subsystem, Domain, Endurance Group, and NVM Set are described in section 3.8. For each namespace, the NVM Set and the Endurance Group [that contain it](#) ~~which~~ the namespace ~~is allocated~~ are reported in the Identify Namespace data structure. The NVM Set to be used for a namespace is based on the value in the NVM Set Identifier field in a create operation. If the NVM Set Identifier field is cleared to 0h in a create operation, then the controller shall choose the NVM Set from which to allocate [capacity to create](#) the namespace.

If the NVM Set Identifier field and the Endurance Group Identifier field are both cleared to 0h in a create operation, then the controller shall choose the Endurance Group and the NVM Set from which to allocate capacity to create the namespace.

If the NVM Set Identifier field is cleared to 0h and the Endurance Group Identifier field is set to a non-zero value in a create operation, then the controller shall choose the NVM Set in the specified Endurance Group from which to allocate capacity to create the namespace.

If the NVM Set Identifier field is set to a non-zero value and the Endurance Group Identifier field is cleared to 0h in a create operation, then the controller shall abort the command with a status code of Invalid Field in Command.

If the NVM Set Identifier field and the Endurance Group Identifier field are both set to non-zero values in a create operation and the specified NVM Set exists in the specified Endurance Group, then the controller shall allocate capacity for the created namespace from the specified NVM Set.

...

Modify a portion of section 8.12.1.1 as shown below:

8.12 Namespace Write Protection

...

8.12.1 Namespace Write Protection – Theory of Operation

...

8.12.1.1 Namespace Write Protection – Command Interactions

...

Figure 431: Commands Allowed when Specifying a Write Protected NSID

Admin Command Set	NVM Command Set
Device Self-test	Compare
Directive Send ¹	Dataset Management ¹
Directive Receive ³	Read
Get Features	Reservation Register
Get Log Page	Reservation Report
Identify	Reservation Acquire
Namespace Attachment	Reservation Release
Security Receive ¹	Vendor Specific ¹
Security Send ¹	Flush ²
Set Features ¹	Verify
Vendor Specific ¹	
NOTES:	
1. The controller shall fail commands if the specified action attempts to modify the non-volatile storage medium of the specified namespace.	
2. A Flush command shall complete successfully with no effect. All volatile write cache data and metadata associated with the specified namespace is written to non-volatile storage medium as part of transitioning to the write protected state (refer to section 5.27.1.28).	
3. A Directive Receive command which attempts to allocate streams resources shall be aborted with a status code of Namespace is Write Protected.	

...

Modify a portion of section 8.13.3 as shown below:

8.13 NVMe over Fabrics Secure Channel and In-band Authentication

...

8.13.3 NVMe In-band Authentication Followed by Secure Channel Establishment

It is possible to leverage an authentication transaction to generate shared key material to use as pre-shared key (PSK) to establish a secure channel (e.g., with IPsec or TLS). In this case, the PSK generated to set up a secure channel on the Admin Queue may be reused to set up additional secure channels on the I/O queues. Figure 436 shows an example of this possibility for TLS.

...

Modify a portion of section 8.13.5.5 as shown below:

8.13.5 DH-HMAC-CHAP Protocol

...

8.13.5.5 DH-HMAC-CHAP_Success1 Message

...

Figure 450: DH-HMAC-CHAP_Success1 message format

Bytes	Description
...	
15+HL:16	Response Value (RVAL): Response R ₂ , if valid (i.e., if the RVALID field is set to 01h), cleared to 0hzero otherwise

...

Modify a portion of section 8.13.5.9 as shown below:

8.13.5.9 Generated PSK for TLS

...

The generated PSK for TLS shall be computed applying the HMAC function using the hash function H() selected by the HashID parameter in the DH-HMAC-CHAP_Challenge message with the session key K_s as key to the concatenation of the two challenges C₁ and C₂ (i.e., generated PSK = HMAC(K_s, C₁ || C₂)). The generated PSK used to set up a TLS secure channel on the Admin Queue may be reused to set up additional TLS secure channels on the I/O queues (refer to the PSK Reuse section of the NVMe TCP Transport Specification). The lifetime of this generated PSK should be no more than ten minutes; this requires authentication for I/O queues created after this time.

...

Modify a portion of section 8.24 as shown below:

8.24 Telemetry

...

The trigger for the collection for host-initiated data is typically a system crash, but may also be initiated during normal operation. The host proceeds with a host-initiated data collection by submitting the Get Log Page command for the Telemetry Host-Initiated log page with the Create Telemetry Host-Initiated Data bit set to '1' in the Log Specific Parameter field. The controller should complete the command quickly (e.g., in less than one second) to avoid a user rebooting the system prior to completion of the data collection.

...

Description of NVM Express NVM Command Set specification 1.0b changes

Modify a portion of section 3.2.1 as shown below:

3 I/O Commands for the NVM Command Set

...

3.2 NVM Command Set Commands

...

3.2.1 Compare command

The Compare command reads the logical blocks specified by the command from the [non-volatile storage](#) medium and compares the data read to a comparison data buffer transferred as part of the command. If the data read from the controller and the comparison data buffer are equivalent with no mismatches, then the command completes successfully. If there is any mismatch, the command completes with an error of Compare Failure.

...

Figure 23: Compare – Command Dword 12

Bits	Description
...	
29:26	Protection Information Field (PRINFO): Specifies the protection information action and check field, as defined in Figure 9. The Protection Information Action (PRACT) bit shall be cleared to '0'. If the Protection Information Check (PRCHK) field is non-zero, a check is performed on the logical block read from NVM (refer to section 5.2.2.4).
...	Reserved

...

3.2.1.1 Command Completion

...

Figure 26: Compare – Command Specific Status Values

Value	Description
81h	Invalid Protection Information: The Protection Information Field (PRINFO) field (refer to Figure 23) settings specified in the command are invalid for the Protection Information with which the namespace was formatted (refer to the PI field in Figure 78 and the DPS field in Figure 97) or the EILBRT field is invalid (refer to section 5.2.2.5).

...

Modify a portion of section 3.2.2 as shown below:

3.2.2 Copy command

...

Figure 30: Copy – Command Dword 12

Bits	Description								
...									
29:26	Protection Information Field Write (PRINFOW): Specifies the protection information action and check field, as defined in Figure 9, to be used for the write portion of the copy operation.								
25	Reserved								
24	Storage Tag Check Write (STCW): This bit specifies the Storage Tag field shall be checked as part of end-to-end data protection processing as defined in Figure 10, to be used for the write portion of the copy operation.								
23:20	Directive Type (DTYPE): Specifies the Directive Type associated with the Directive Specific field (refer to the Directives section in the NVMe Base Specification) used for the write portion of the copy operation.								
19:16	Reserved								
15:12	Protection Information Field Read (PRINFOR): Specifies the protection information action and check field, as defined in Figure 9, to be used for the read portion of the copy operation specified by each Source Range Entries.								
11:08	Descriptor Format: Specifies the type of the Copy Descriptor Format that is used. The Copy Descriptor Format specifies the starting LBA, number of logical blocks, and parameters associated with the read portion of the operation.								
	<table><tr><th>Copy Descriptor Format Type</th><th>Description</th></tr><tr><td>0h</td><td>Source Range Entries Descriptor Format 0h is used (refer to Figure 34).</td></tr><tr><td>1h</td><td>Source Range Entries Descriptor Format 1h is used (refer to Figure 35).</td></tr><tr><td>All Others</td><td>Reserved</td></tr></table>	Copy Descriptor Format Type	Description	0h	Source Range Entries Descriptor Format 0h is used (refer to Figure 34).	1h	Source Range Entries Descriptor Format 1h is used (refer to Figure 35).	All Others	Reserved
	Copy Descriptor Format Type	Description							
	0h	Source Range Entries Descriptor Format 0h is used (refer to Figure 34).							
	1h	Source Range Entries Descriptor Format 1h is used (refer to Figure 35).							
All Others	Reserved								
07:00	Number of Ranges (NR): Specifies the number of Source Range Entries that are specified in the command. This is a 0's-based value.								

...

3.2.2.1 Command Completion

...

Figure 37: Copy – Command Specific Status Values

Value	Description
81h	Invalid Protection Information: The protection information specified by the command is invalid due to: <ul style="list-style-type: none"> The Protection Information Field-Read (PRINFOR) field or Protection Information Field Write (PRINFOW) field (refer to Figure 30) containing an invalid value for the Protection Information with which the namespace was formatted (refer to the PI field in the Format NVM Command section in the NVMe Base Specification and the DPS field in Figure 97) the ILBRT field being invalid (refer to section 5.2.2.5); or the EILBRT field in a Source Range Entry being invalid (refer to section 5.2.2.5). <add a period>
...	

...

Modify a portion of section 3.2.4 as shown below:

3.2.4 Read command

...

Figure 48: Read – Command Dword 12

Bits	Description
...	
29:26	Protection Information Field (PRINFO): Specifies the protection information action and check field, as defined in Figure 9.
...	

...

3.2.4.1 Command Completion

...

Figure 52: Read – Command Specific Status Values

Value	Description
...	
81h	Invalid Protection Information: The Protection Information Field (PRINFO) field (refer to Figure 48) settings specified in the command are invalid for the Protection Information with which the namespace was formatted (refer to the PI field in Figure 78 and the DPS field in Figure 97) or the EILBRT field is invalid (refer to section 5.2.2.5).

...

Modify a portion of section 3.2.5 as shown below:

3.2.5 Verify command

...

Figure 55: Verify – Command Dword 12

Bits	Description
...	
29:26	Protection Information Field (PRINFO): Specifies the protection information action and check field, as defined in Figure 9. The Protection Information Check (PRCHK) field in the PRINFO field specifies the protection information to be checked by the Verify operation. The Protection Information Action (PRACT) bit in the PRINFO field is cleared to '0' by the host. If the PRACT bit is not cleared to '0', then the controller shall abort the command with a status of Invalid Field in Command.
...	

...

3.2.5.1 Command Completion

...

Figure 58: Verify – Command Specific Status Values

Value	Description
81h	Invalid Protection Information: The Protection Information Field (PRINFO) field (refer to Figure 55) settings specified in the command are invalid for the Protection Information with which the namespace was formatted (refer to the PI field in Figure 78 and the DPS field in Figure 97) or the EILBRT field is invalid (refer to section 5.2.2.5).

...

Modify a portion of section 3.2.6 as shown below:

3.2.6 Write command

...

Figure 63: Write – Command Dword 12

Bits	Description
...	
29:26	Protection Information Field (PRINFO): Specifies the protection information action and check field, as defined in Figure 9.
...	

...

3.2.6.1 Command Completion

...

Figure 67: Write – Command Specific Status Values

1	Value	2	Description
...			
	81h		Invalid Protection Information: The Protection Information Field (PRINFO) field (refer to Figure 63) settings specified in the command are invalid for the Protection Information with which the namespace was formatted (refer to the PI field in Figure 78 and the DPS field in Figure 97) or the ILBRT field is invalid (refer to section 5.2.2.5).
...			

...

Modify a portion of section 3.2.8 as shown below:

3.2.8 Write Zeroes command

...

Figure 73: Write Zeroes – Command Dword 12

3	Bits	4	Description
...			
	29:26		Protection Information Field (PRINFO): Specifies the protection information action and check field, as defined in Figure 9. The Protection Information Check (PRCHK) field shall be cleared to 000b.
...			

...

3.2.8.1 Command Completion

...

Figure 76: Write Zeroes – Command Specific Status Values

5	Value	6	Description
...			
	81h	Invalid Protection Information: The Protection Information Field -(PRINFO) field (refer to Figure 73) settings specified in the command are invalid for the Protection Information with which the namespace was formatted (refer to the PI field in Figure 78 and the DPS field in Figure 97) or the ILBRT field is invalid (refer to section 5.2.2.5).	
...			

...

Modify a portion of section 4.1.4.1 as shown below:

4 Admin Commands for the NVM Command Set

...

4.1 Admin Command behavior for the NVM Command Set

...

4.1.4 Get Log Page command

...

4.1.4.1 Error Information (Log Identifier 01h)

The Error Information ~~L~~og ~~P~~age is as defined in the NVMe Base Specification. Figure 90 describes the NVM Command Set specific definition of the User Data field.

...

Modify a portion of section 4.1.4.2 as shown below:

4.1.4.2 SMART / Health Information (02h)

The SMART / Health Information ~~L~~og ~~P~~age is as defined in the NVMe Base Specification. For the Data Units Read and Data Units Written fields, when the LBA size is a value other than 512 bytes, the controller shall convert the amount of data read to 512 byte units.

...

Modify a portion of section 4.1.4.3 as shown below:

4.1.4.3 Device Self-test (Log Identifier 06h)

The Device Self-test ~~L~~og ~~P~~age is as defined in the NVMe Base Specification. Figure 91 describes the NVM Command Set specific definition of the Failing User Data Reference field.

Modify a portion of section 4.1.4.4 as shown below:

4.1.4.4 Persistent Event (Log Identifier 0Dh)

The Persistent Event ~~Log~~ ~~P~~age is as defined in the NVMe Base Specification. Figure 92 describes the NVM Command Set specific definition of the I/O Command Set specific fields within the Change Namespace Event Data Format (Event Type 06h) (refer to the NVMe Base Specification).

...

Modify a portion of section 4.1.5.1 as shown below:

4.1.5 Identify Command

...

4.1.5.1. NVM Command Set Identify Namespace Data Structure (CNS 00h)

...

Figure 97: Identify – Identify Namespace Data Structure, NVM Command Set

Bytes	O/M ¹	Description
...		
25	M	Number of LBA Formats (NLBAF): This field defines the number of supported LBA data size and metadata size combinations supported by the namespace. LBA formats shall be allocated in order (starting with 0) and packed sequentially starting at the LBA Format 0 Support (LBAF0) field. This is a 0's based value. The maximum number of LBA formats that may be indicated as supported is: a) 16 if the LBA Format Extension Enable (LBAFEE) field is cleared to 0h in the Host Behavior Support feature (refer to the Host Behavior Support section in the NVMe Base Specification); or b) 64 if the LBAFEE field is set to 1h in the Host Behavior Support feature (refer to the Host Behavior Support section in the NVMe Base Specification).
...		

...

Modify a portion of section 5.2.2.1 as shown below:

5 Extended Capabilities

5.2 End-to-end Data Protection

...

5.2.2 PRACT Bit

...

5.2.2.1 Protection Information and Write Commands

...

If the namespace is formatted with protection information and the PRACT bit is set to '1', then:

1. If the namespace is formatted with Metadata Size (refer to Figure 98) equal to protection information size (refer to section 5.2.1), then the logical block data is transferred from the host buffer to the controller. As the logical block data passes through the controller, the controller generates and appends protection information to the end of the logical block data, and the logical block data and protection information are written to NVM (i.e., the metadata is not resident within the host buffer); and
2. If the namespace is formatted with Metadata Size greater than protection information size, then the logical block data and the metadata are transferred from the host buffer to the controller. As the metadata passes through the controller, the controller overwrites the protection information portion of the metadata **without checking the protection information portion regardless of PRCHK settings**. The logical block data and metadata are written to the NVM (i.e., the metadata field remains the same size in the NVM and the host buffer). The location of the protection information within the metadata is configured when the namespace is formatted (refer to the DPS field in Figure 97).

...

Modify a portion of section 5.3 as shown below:

5.3 Namespace Management

Namespace Management **capability** operates as defined in the NVMe Base Specification with additional capabilities specifically for the NVM Command Set.

The NVM Command Set supports reporting of Namespace Granularity as I/O Command Set specific Namespace Management **capability** content. The Namespace Granularity List defined in Figure 103 is requested by host software using the Identify command with CNS set to 16h.

If the controller supports reporting of Namespace Granularity, then the Namespace Granularity Descriptor List (refer to Figure 103) contains one or more Namespace Granularity Descriptors (refer to Figure 104) indicating the size granularity and the capacity granularity **with at** which the controller **creates allocates** namespaces.

The size granularity and the capacity granularity are hints which may be used by the host to minimize the capacity that is allocated for a namespace and that is not able to be addressed by logical block addresses. The granularities are used in specifying values for the Namespace Size (NSZE) **field** and Namespace Capacity (NCAP) **fields** of the data structure used for the create operation of the Namespace Management command (refer to **Figure 105 the Namespace Management command section in the NVMe Base Specification**).

...

Modify a portion of section 5.8.2.1 as shown below:

5.8 Command Set Specific Capability

...

5.8.2 Improving Performance through I/O Size and Alignment Adherence

...

5.8.2.1 Improved I/O examples (non-normative)

...

Namespace Optimal Write Size (NOWS) (refer to Figure 97) is intended to supplement NVM Sets Optimal Write Size as NOWS provides a mechanism to report the optimal write size that scales to a multiple namespace per NVM Set use case, but also covers the use case where there is a single namespace ~~that exists allocated~~ in an NVM Set. Namespaces should report NOWS as a multiple of NPWG. When constructing write operations, the host should minimally construct writes that meet the recommendations of NPWG and NPWA, but may achieve optimal write performance by constructing writes that meet the recommendation of NOWS.

...

Description of NVM Express Zoned Namespace Command Set specification 1.1b changes

Modify a portion of section 3.4.1 as shown below:

3 I/O Commands for the Zoned Namespace Command Set

...

3.4 Zoned Namespace Command Set I/O Commands

...

3.4.1 Zone Append command

...

Figure 26: Zone Append – Command Dword 12

Bits	Description
...	
29:26	Protection Information Field (PRINFO): Specifies the Protection Information field, as defined in the NVM Command Set Specification.
...	

...

3.4.1.2 Command Completion

...

Figure 30: Zone Append – Command Specific Status Values

Value	Description
81h	Invalid Protection Information: The Protection Information (PRINFO) field settings specified in the command are invalid for the Protection Information with which the namespace was formatted or the ILBRT field is invalid (refer to section 3.4.1).
B8h	Zone Boundary Error: The command specifies logical blocks in more than one zone.
B9h	Zone Is Full: The accessed zone is in the ZSF:Full state.
BAh	Zone Is Read Only: The accessed zone is in the ZSRO:Read Only state.
BBh	Zone Is Offline: The accessed zone is in the ZSO:Offline state.
BDh	Too Many Active Zones: The controller does not allow additional active zones.
BEh	Too Many Open Zones: The controller does not allow additional open zones.

...

Modify a portion of section 3.4.2 as shown below:

3.4.2 Zone Management Receive command

...

Figure 34: Zone Management Receive – Command Dword 13

Bits	Description																																			
...																																				
15:08	Zone Receive Action Specific-Field:																																			
	<table><tr><th>Zone Receive Action</th><th>Description</th></tr><tr><td rowspan="2">Report Zones</td><td>Reporting Options:</td></tr><tr><td><table><tr><th>Value</th><th>Description</th></tr><tr><td>0h</td><td>List all zones.</td></tr><tr><td>1h</td><td>List the zones in the ZSE:Empty state.</td></tr><tr><td>2h</td><td>List the zones in the ZSIO:Implicitly Opened state.</td></tr><tr><td>3h</td><td>List the zones in the ZSEO:Explicitly Opened state.</td></tr><tr><td>4h</td><td>List the zones in the ZSC:Closed state.</td></tr><tr><td>5h</td><td>List the zones in the ZSF:Full state.</td></tr><tr><td>6h</td><td>List the zones in the ZSRO:Read Only state.</td></tr><tr><td>7h</td><td>List the zones in the ZSO:Offline state.</td></tr><tr><td>8h</td><td>Reserved</td></tr><tr><td>9h</td><td>List all zones that have the zone attribute in the Zone Attribute field of the Zone Descriptor data structure (refer to Figure 37):<ul style="list-style-type: none">Reset Zone Recommended bit set to '1';Finish Zone Recommended bit set to '1'; orZone Finished by Controller bit set to '1'.</td></tr><tr><td>10h</td><td rowspan="3">Reserved</td></tr><tr><td>11h</td></tr><tr><td>3Fh</td></tr><tr><td>All other values</td><td>Reserved</td></tr></table></td></tr><tr><td>All other values</td><td>Reserved</td></tr></table>	Zone Receive Action	Description	Report Zones	Reporting Options:	<table><tr><th>Value</th><th>Description</th></tr><tr><td>0h</td><td>List all zones.</td></tr><tr><td>1h</td><td>List the zones in the ZSE:Empty state.</td></tr><tr><td>2h</td><td>List the zones in the ZSIO:Implicitly Opened state.</td></tr><tr><td>3h</td><td>List the zones in the ZSEO:Explicitly Opened state.</td></tr><tr><td>4h</td><td>List the zones in the ZSC:Closed state.</td></tr><tr><td>5h</td><td>List the zones in the ZSF:Full state.</td></tr><tr><td>6h</td><td>List the zones in the ZSRO:Read Only state.</td></tr><tr><td>7h</td><td>List the zones in the ZSO:Offline state.</td></tr><tr><td>8h</td><td>Reserved</td></tr><tr><td>9h</td><td>List all zones that have the zone attribute in the Zone Attribute field of the Zone Descriptor data structure (refer to Figure 37):<ul style="list-style-type: none">Reset Zone Recommended bit set to '1';Finish Zone Recommended bit set to '1'; orZone Finished by Controller bit set to '1'.</td></tr><tr><td>10h</td><td rowspan="3">Reserved</td></tr><tr><td>11h</td></tr><tr><td>3Fh</td></tr><tr><td>All other values</td><td>Reserved</td></tr></table>	Value	Description	0h	List all zones.	1h	List the zones in the ZSE:Empty state.	2h	List the zones in the ZSIO:Implicitly Opened state.	3h	List the zones in the ZSEO:Explicitly Opened state.	4h	List the zones in the ZSC:Closed state.	5h	List the zones in the ZSF:Full state.	6h	List the zones in the ZSRO:Read Only state.	7h	List the zones in the ZSO:Offline state.	8h	Reserved	9h	List all zones that have the zone attribute in the Zone Attribute field of the Zone Descriptor data structure (refer to Figure 37): <ul style="list-style-type: none">Reset Zone Recommended bit set to '1';Finish Zone Recommended bit set to '1'; orZone Finished by Controller bit set to '1'.	10h	Reserved	11h	3Fh	All other values	Reserved	All other values	Reserved
	Zone Receive Action	Description																																		
	Report Zones	Reporting Options:																																		
		<table><tr><th>Value</th><th>Description</th></tr><tr><td>0h</td><td>List all zones.</td></tr><tr><td>1h</td><td>List the zones in the ZSE:Empty state.</td></tr><tr><td>2h</td><td>List the zones in the ZSIO:Implicitly Opened state.</td></tr><tr><td>3h</td><td>List the zones in the ZSEO:Explicitly Opened state.</td></tr><tr><td>4h</td><td>List the zones in the ZSC:Closed state.</td></tr><tr><td>5h</td><td>List the zones in the ZSF:Full state.</td></tr><tr><td>6h</td><td>List the zones in the ZSRO:Read Only state.</td></tr><tr><td>7h</td><td>List the zones in the ZSO:Offline state.</td></tr><tr><td>8h</td><td>Reserved</td></tr><tr><td>9h</td><td>List all zones that have the zone attribute in the Zone Attribute field of the Zone Descriptor data structure (refer to Figure 37):<ul style="list-style-type: none">Reset Zone Recommended bit set to '1';Finish Zone Recommended bit set to '1'; orZone Finished by Controller bit set to '1'.</td></tr><tr><td>10h</td><td rowspan="3">Reserved</td></tr><tr><td>11h</td></tr><tr><td>3Fh</td></tr><tr><td>All other values</td><td>Reserved</td></tr></table>	Value	Description	0h	List all zones.	1h	List the zones in the ZSE:Empty state.	2h	List the zones in the ZSIO:Implicitly Opened state.	3h	List the zones in the ZSEO:Explicitly Opened state.	4h	List the zones in the ZSC:Closed state.	5h	List the zones in the ZSF:Full state.	6h	List the zones in the ZSRO:Read Only state.	7h	List the zones in the ZSO:Offline state.	8h	Reserved	9h	List all zones that have the zone attribute in the Zone Attribute field of the Zone Descriptor data structure (refer to Figure 37): <ul style="list-style-type: none">Reset Zone Recommended bit set to '1';Finish Zone Recommended bit set to '1'; orZone Finished by Controller bit set to '1'.	10h	Reserved	11h	3Fh	All other values	Reserved						
	Value	Description																																		
	0h	List all zones.																																		
	1h	List the zones in the ZSE:Empty state.																																		
	2h	List the zones in the ZSIO:Implicitly Opened state.																																		
	3h	List the zones in the ZSEO:Explicitly Opened state.																																		
	4h	List the zones in the ZSC:Closed state.																																		
	5h	List the zones in the ZSF:Full state.																																		
	6h	List the zones in the ZSRO:Read Only state.																																		
7h	List the zones in the ZSO:Offline state.																																			
8h	Reserved																																			
9h	List all zones that have the zone attribute in the Zone Attribute field of the Zone Descriptor data structure (refer to Figure 37): <ul style="list-style-type: none">Reset Zone Recommended bit set to '1';Finish Zone Recommended bit set to '1'; orZone Finished by Controller bit set to '1'.																																			
10h	Reserved																																			
11h																																				
3Fh																																				
All other values	Reserved																																			
All other values	Reserved																																			
...																																				

...

Modify a portion of section 4.1.1 as shown below:

4 Admin Commands for the Zoned Namespace Command Set

...

4.1 Admin Command behavior for the Zoned Namespace Command Set

...

4.1.1 Asynchronous Event Request Command

...

Figure 42: Asynchronous Event Information – Notice, Zoned Namespace Command Set

Value	Description
EFh	<p>Zone Descriptor Changed: The Zone Descriptor data structure for a zone in a specific zoned namespace (refer to Figure 43) has changed. The Zone Descriptor of the zone that changed is indicated in the Changed Zone List log page. To clear this event, host software reads the Zone Changed List log page using the Get Log Page command with the Retain Asynchronous Event bit cleared to '0'.</p> <p>For a specific zone, a Zone Descriptor data structure change that occurred, <i><added comma></i> for any of the following reasons, shall not generate a Zone Descriptor Changed event and shall not cause modifications to the Changed Zones List log page:</p> <ul style="list-style-type: none"> a) a Zone Management Send command that specified that zone; b) a Zone Management Send command that specified all zones; c) a write operation that transitioned that zone: <ul style="list-style-type: none"> i. from the ZSE:Empty state to the ZSIO:Implicitly Opened state; ii. from the ZSIO:Implicitly Opened state to the ZSF:Full state; iii. from the ZSEO:Explicitly Opened state to the ZSF:Full state; or iv. from the ZSC:Closed state to the ZSIO:Implicitly Opened state; d) the controller transitioning that zone to the ZSF:Full state due to an NVM Subsystem Reset; and e) the controller transitioning that zone to the ZSC:Closed state (refer to section 2.1.1.4.1).

...

Modify a portion of section 4.1.5.1 as shown below:

4.1.5 Identify Command

...

4.1.5.1 I/O Command Set Specific Identify Namespace Data Structure for the Zoned Namespace Command Set (CNS 05h, CSI 02h)

...

Figure 48: I/O Command Set Specific Identify Namespace Data Structure for the Zoned Namespace Command Set

Bytes	O/M ¹	Description
...		
LBA Format Extensions		
2831:2816	M	LBA Format 0 Extension (LBAFE0): This field indicates the LBA format Extension 0 that is supported by the controller. The LBA Format Extension Zone format field is defined in Figure 49.
2847:2832	O	LBA Format 1 Extension (LBAFE1): This field indicates the LBA format Extension 1 that is supported by the controller. The LBA Format Extension field is defined in Figure 49.
...		

Figure 48: I/O Command Set Specific Identify Namespace Data Structure for the Zoned Namespace Command Set

Bytes	O/M ¹	Description
3839:3824	O	LBA Format 63 Extension (LBAFE63): This field indicates the LBA format Extension 63 that is supported by the controller. The LBA Format Extension field is defined in Figure 49.
4095:3840	O	Vendor Specific
NOTES: 1. O/M definition: O = Optional, M = Mandatory.		

...

Description of NVM Express Key Value Command Set specification
1.0b changes

Modify a portion of section 4.1.5.1 as shown below:

4 Admin Commands for the Key Value Command Set

...

4.1 Admin Command behavior for the Key Value Command Set

...

4.1.5 Identify Command

...

4.1.5.1 I/O Command Set specific Identify Namespace data structure (CNS 05h, CSI 01h)

Figure 36: Identify – I/O Command Set Specific Identify Namespace Data Structure, Key Value Type Specific

Bytes	O/M ¹	Description
...		
25	M	<p>Number of KV Formats (NKVF): This field defines the number of KV format descriptors supported by the namespace. KV formats shall be allocated in order (starting with 0) and packed sequentially starting at the KV Format 0 Support (KVF0) field. This is a 0's based value.</p> <p>The maximum number of KV formats that may be indicated as supported is 16. The supported KV formats are indicated in bytes 72 to 327 in this data structure.</p> <p>The KV Format fields with a Format Index beyond the value set in this field are invalid and not supported. KV Formats that are valid, but not currently available may be indicated by clearing the KV Key Max Length field to 0h and clearing the KV Value Max Length field to 0h for that KV Format.</p>
...		

...

Description of NVM Express Management Interface 1.2b changes

Modify a portion of section 3.1.1.1 as shown below:

3 Message Transport

...

3.1 NVMe-MI Messages

...

3.1.1 Message Fields

...

3.1.1.1 Message Integrity Check

If the Integrity Check (IC) bit is set to ‘1’, then the Message Integrity Check field contains a 32-bit CRC computed over the contents of the NVMe-MI Message. The 32-bit CRC required by this specification is CRC-32C (Castagnoli) which uses the generator polynomial 1EDC6F41h. The Message Integrity Check is calculated using the following Rocksoft™ Model CRC Algorithm parameters defined in Figure X.

Figure X: Rocksoft™ Model CRC Algorithm parameters

Parameter		Value
Name	÷	"CRC-32C"
Width	÷	32
Poly	÷	1EDC6F41h
Init	÷	FFFFFFFFh
RefIn	÷	True
RefOut	÷	True
XorOut	÷	FFFFFFFFh
Check	÷	E3069283h

...

Modify a portion of section 3.2.2 as shown below:

3.2 Out-of-Band Message Transport

...

3.2.2 Out-of-Band Error Handling

The Management Endpoint shall drop (silently discard) packets for error conditions as specified in the MCTP Base Specification. Some example conditions which result in discarding packets include unexpected middle or end packets. Silently discarded packets also cause the corresponding bit in the Get State Control Primitive Success Response Fields to be set to ‘1’ (refer to Figure 42).

...

Modify a portion of section 4.3.1.1 as shown below:

4 Message Servicing Model

...

4.3 In-Band Tunneling Message Servicing Model

...

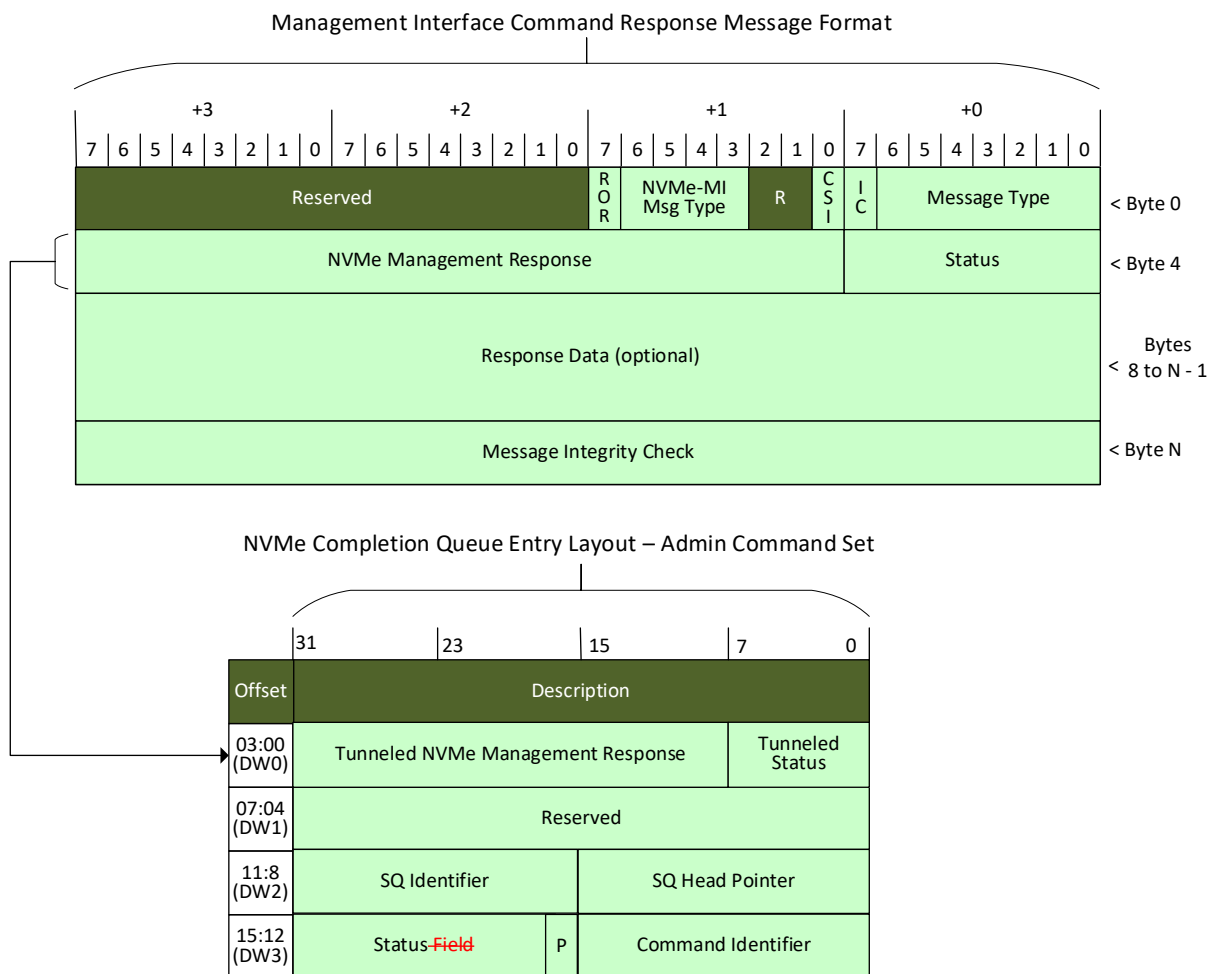
4.3.1 NVMe-MI Send Command

...

4.3.1.1 NVMe-MI Send Command Request Message to NVMe Admin Command SQE Mapping

...

Figure 47: NVMe-MI Send Command Response Message to NVMe Admin Command CQE Mapping Diagram



...

Modify a portion of section 4.3.1.2 as shown below:

4.3.1.2 NVMe-MI Send Command Servicing Model

...

This section describes the NVMe-MI Send command servicing model starting at NVMe Processing as shown in phase 1 of Figure 50. In phase 1, CDW0 to CDW9 are checked for errors per the NVM Express Base Specification. If any errors are encountered in CDW0 to CDW9, then the NVMe-MI Send command is completed with an error status code in the Status **F**field as per the NVM Express Base Specification and the Tunneled Status and Tunneled NVMe Management Response fields shall be cleared to 0h.

If there are no errors in CDW0 to CDW9, then command servicing enters phase 2 where the portion of the tunneled NVMe-MI Command in CDW10 to CDW15 is checked for errors. Note that if there is no Request Data, then CDW10 to CDW15 contain the entire tunneled NVMe-MI Command. If any errors are encountered in the portion of the tunneled NVMe-MI Command in CDW10 to CDW15, then the NVMe-MI Send command is completed with a status code of Successful Completion in the Status **F**field as defined in the NVM Express Base Specification. The Tunneled Status field contains the error Response Message Status for the portion of the tunneled NVMe-MI Command in CDW10 to CDW15 and the Tunneled NVMe Management Response field contains the NVMe Management Response field from the NVMe-MI Command that is being tunneled in-band.

If there are no errors in phase 2, then command servicing enters phase 3 where there is a check to determine if there is any Request Data for the tunneled NVMe-MI Command. If there is no Request Data for the tunneled NVMe-MI Command, then command servicing skips to phase 5. If there is Request Data, then the Request Data is transferred from the buffer pointed to by DPTR. If any errors are encountered transferring the Request Data, then the command is completed with an error status code in the Status **F**field as per the NVM Express Base Specification and the Tunneled Status and Tunneled NVMe Management Response fields shall be cleared to 0h.

If there are no errors transferring the data, then command servicing enters phase 4 where the whole tunneled NVMe-MI Command is constructed from CDW10 to CDW15 and the Request Data that was transferred. If any errors are encountered in the tunneled NVMe-MI Command, then the NVMe-MI Send command is completed with a status code of Successful Completion in the Status **F**field as defined in the NVM Express Base Specification. The Tunneled Status field contains the appropriate error Response Message Status and the Tunneled NVMe Management Response field contains the NVMe Management Response field from the NVMe-MI Command that is being tunneled in-band.

If there are no errors in phase 4, then command servicing enters phase 5 where the tunneled NVMe-MI Command finishes processing. If any errors are encountered processing the tunneled NVMe-MI Command, then the NVMe-MI Send command is completed with a status code of Successful Completion in the Status **F**field as defined in the NVM Express Base Specification and the Tunneled Status field contains the appropriate error Response Message Status. If the tunneled NVMe-MI Command is processed successfully, then the NVMe-MI Send command is completed with a status code of Successful Completion in the Status **F**field as defined in the NVM Express Base Specification. The Tunneled Status field contains a Response Message Status of Success for the tunneled NVMe-MI Command and the Tunneled NVMe Management Response field contains the NVMe Management Response field from the NVMe-MI Command that is being tunneled in-band.

...

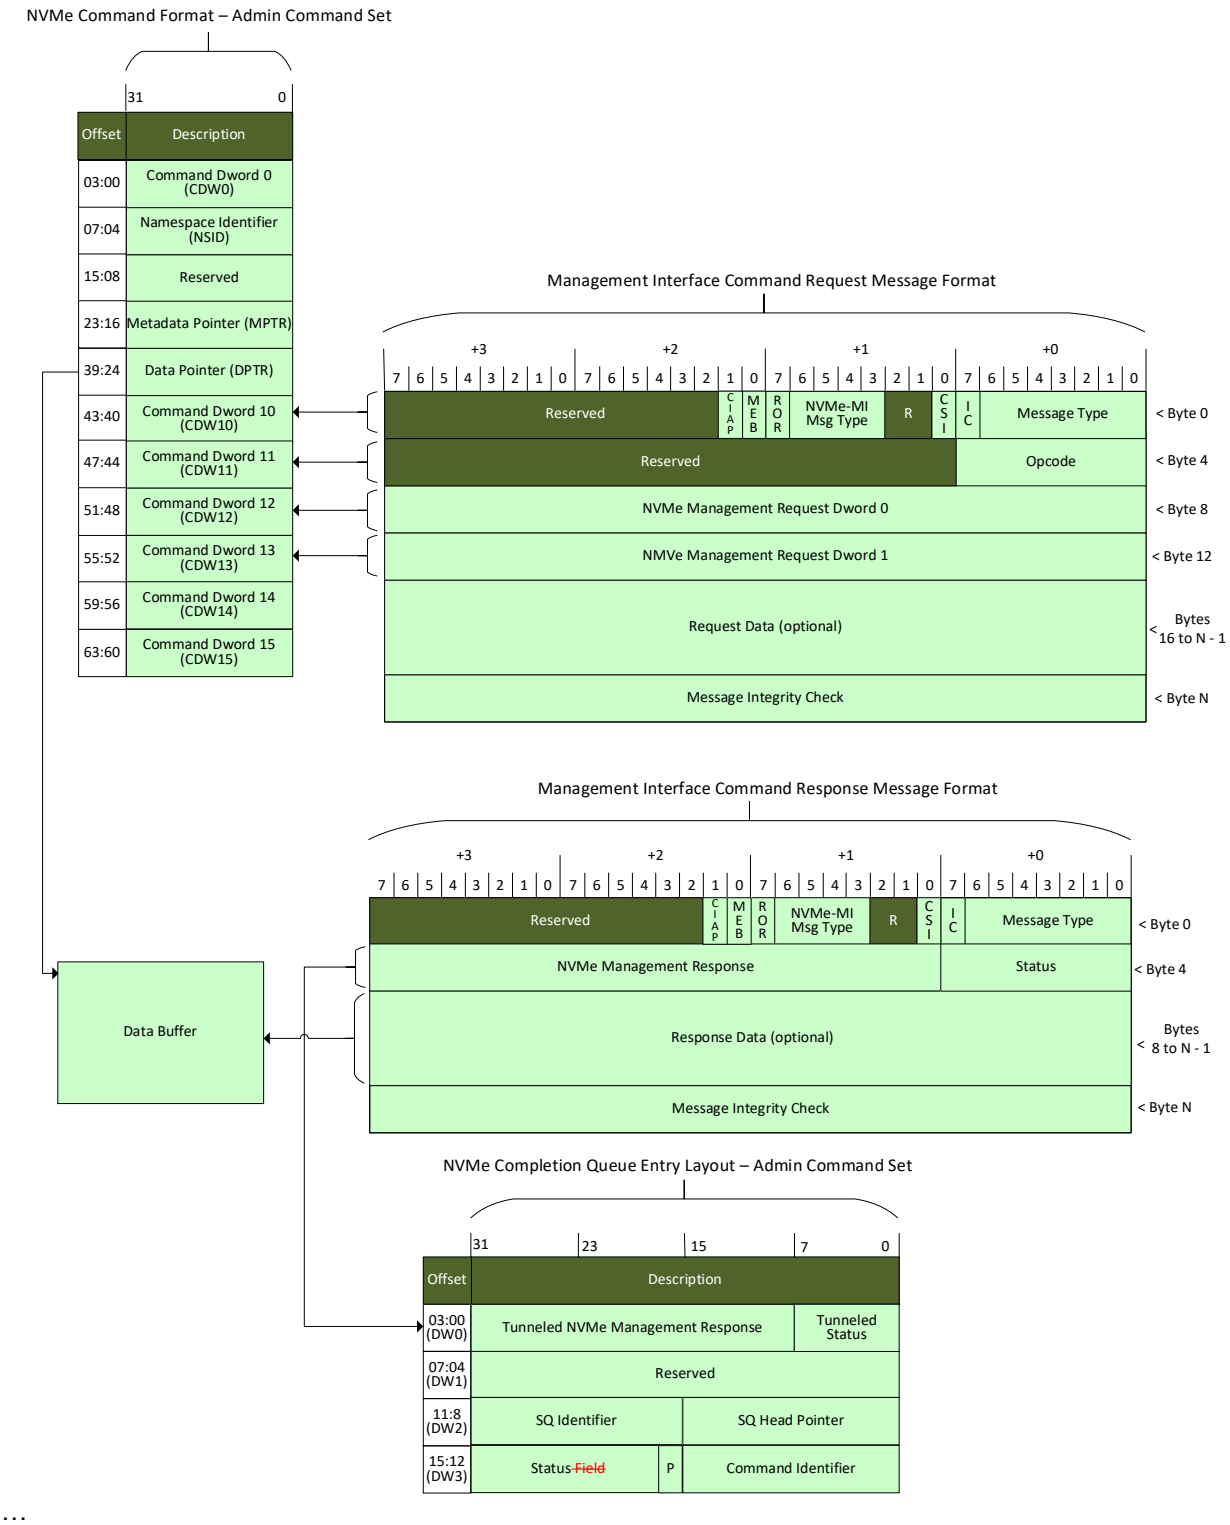
Modify a portion of section 4.3.2.1 as shown below:

4.3.2 NVMe-MI Receive Command

...

4.3.2.1 NVMe-MI Receive Command Request Message to NVMe Admin Command SQE Mapping

Figure 51: NVMe-MI Receive Command Request/Response Message to NVMe Admin Command SQE/CQE Mapping Diagram



Modify a portion of section 4.3.2.2 as shown below:

4.3.2.2 NVMe-MI Receive Command Servicing Model

...

This section describes the NVMe-MI Receive command servicing model starting at NVMe Processing as shown in phase 1 of Figure 55. In phase 1, CDW0 to CDW9 are checked for errors per the NVM Express Base Specification. If any errors are encountered in CDW0 to CDW9, then the command is completed with an error status code in the Status **F**field as per the NVM Express Base Specification and the Tunneled Status and Tunneled NVMe Management Response fields shall be cleared to 0h.

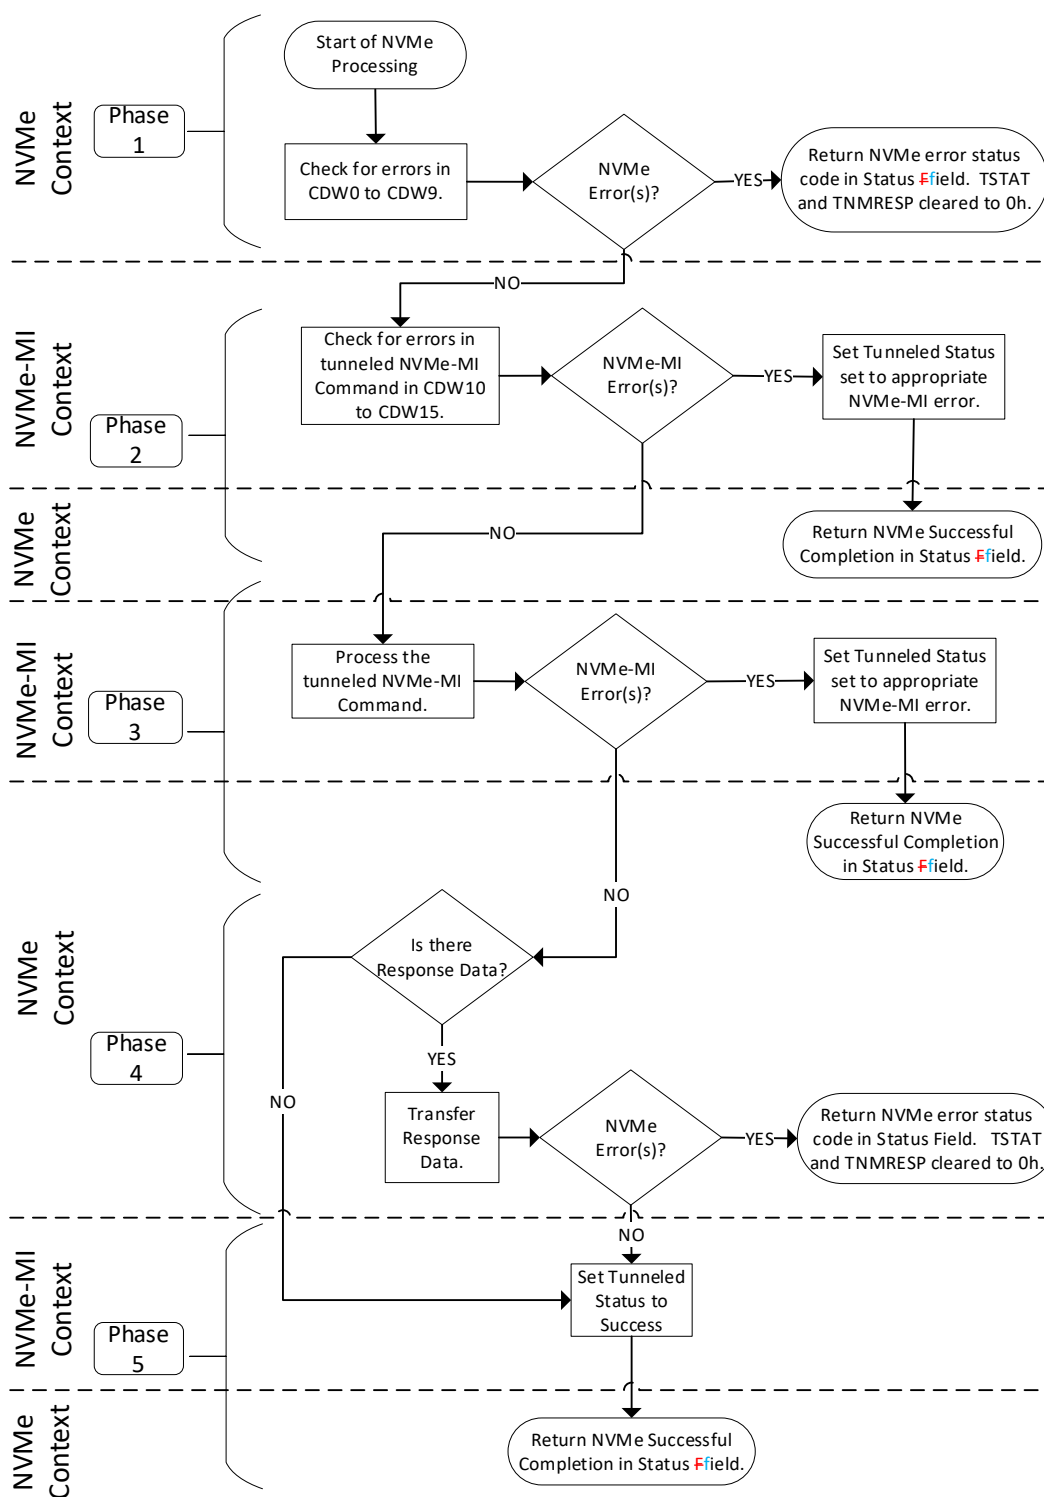
If there are no errors in CDW0 to CDW9, then command servicing enters phase 2 where the tunneled NVMe-MI Command in CDW10 to CDW15 is checked for errors. If any errors are encountered in the tunneled NVMe-MI Command in CDW10 to CDW15, then the NVMe-MI Receive command is completed with a status code of Successful Completion in the Status **F**field as defined in the NVM Express Base Specification. The Tunneled Status field contains the appropriate error Response Message Status and the Tunneled NVMe Management Response field contains the NVMe Management Response field from the NVMe-MI Command that is being tunneled in-band.

If there are no errors in phase 2, then command servicing enters phase 3 where the tunneled NVMe-MI Command finishes processing. If any errors are encountered processing the tunneled NVMe-MI Command, then the NVMe-MI Receive command is completed with a status code of Successful Completion in the Status **F**field as defined in the NVM Express Base Specification. The Tunneled Status field contains the appropriate error Response Message Status and the Tunneled NVMe Management Response field contains the NVMe Management Response field from the NVMe-MI Command that is being tunneled in-band.

If there are no errors in phase 3, then command servicing enters phase 4 where there is a check to determine if there is any Response Data for the tunneled NVMe-MI Command. If there is no Response Data for the tunneled NVMe-MI Command, then command servicing skips to phase 5. If there is Response Data, then the Response Data is transferred to the buffer pointed to by DPTR. If any errors are encountered transferring the Response Data then the command is completed with an error status code in the Status **F**field as per the NVM Express Base Specification and the Tunneled Status and Tunneled NVMe Management Response fields shall be cleared to 0h.

If there are no errors in phase 4, then command servicing enters phase 5 where the NVMe-MI Receive command is completed with a status code of Successful Completion in the Status **F**field as defined in the NVM Express Base Specification. The Tunneled Status field contains a Response Message Status of Success for the tunneled NVMe-MI Command and the Tunneled NVMe Management Response field contains the NVMe Management Response field from the NVMe-MI Command that is being tunneled in-band.

Figure 55: NVMe-MI Receive Command Servicing Model



...

Modify a portion of section 5.7 as shown below:

5 Management Interface Command Set

...

5.7 Read NVMe-MI Data Structure

...

Figure 96: PCIe Port Specific Data

Bytes	Description																
08	<p>PCIe Maximum Payload Size: This field indicates the Max_Payload_SizeMax_Payload_Size setting for the specified PCIe port (refer to the PCI Express Base Specification). If the link is not active, this field should be cleared to 0h.</p> <table><tr><th>Value</th><th>Definition</th></tr><tr><td>0h</td><td>128 bytes</td></tr><tr><td>1h</td><td>256 bytes</td></tr><tr><td>2h</td><td>512 bytes</td></tr><tr><td>3h</td><td>1 KiB</td></tr><tr><td>4h</td><td>2 KiB</td></tr><tr><td>5h</td><td>4 KiB</td></tr><tr><td>6h to FFh</td><td>Reserved</td></tr></table> <p>The value reported in this field by ARI Devices and Non-ARI Multi-Function Devices (refer to the PCI Express Base Specification) whose Max Payload Size settings are identical across all Functions is the setting in Function 0. The value reported in this field by non-ARI Multi-Function Devices whose Max Payload Size settings are not identical across all Functions is implementation specific.</p>	Value	Definition	0h	128 bytes	1h	256 bytes	2h	512 bytes	3h	1 KiB	4h	2 KiB	5h	4 KiB	6h to FFh	Reserved
Value	Definition																
0h	128 bytes																
1h	256 bytes																
2h	512 bytes																
3h	1 KiB																
4h	2 KiB																
5h	4 KiB																
6h to FFh	Reserved																
...																	

...

Modify a portion of section 6.2 as shown below:

6 NVM Express Admin Command Set

...

6.2 Status

A Response Message for an NVMe Admin Command may contain two status fields. The first status field, contained in Byte 4 of the Response Message, is defined by this specification, and the second Status ~~F~~field, if present, is contained in Completion Queue Entry Dword 3 and defined in the NVM Express Base Specification.

An NVMe Admin Command Request Message is well formed if it does not contain any of the following errors:

- Invalid Opcode (e.g., the opcode is not listed in Figure 115);
- Invalid Parameter (e.g., the Controller ID field specifies a Controller ID not implemented in the NVM Subsystem);
- Invalid Command Size (e.g., the Request Message does not contain a complete command); or

- Invalid Command Input Data Size (e.g., the Request Data field is larger than the size specified in the Data Length field).

If the NVMe Admin Command Request Message is well formed, then a Success Response is transmitted. The Success Response contains the status associated with NVMe Admin Command in the Status ~~F~~field of Completion Queue Entry Dword 3. The Status ~~F~~field contains any NVM Express Base and I/O Command Set specifications specific status codes (e.g., Success or Invalid Parameter).

...

Modify a portion of section 6.3 as shown below:

6.3 Get Log Page

Figure 122 defines the log pages that are mandatory, optional, and prohibited for SMBus/I2C and PCIe VDM Management Endpoint on NVMe Storage Devices and NVMe Enclosures. The set of optional log pages supported on each Management Endpoint ~~is are~~ allowed to differ (refer to NVM Express Base Specification).

Figure 122: 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
...			
Sanitize Status	81h	O	O
Changed Zone List ⁵	BFh	O	O
Vendor Specific	C0h to FFh	O	O
...			

...

6.5 Set Features and Get Features

...

Figure 126 define the features that are mandatory, optional, and prohibited for SMBus/I2C and PCIe VDM Management Endpoints on NVMe Storage Devices and NVMe Enclosures. The set of optional features supported on each Management Endpoint ~~is are~~ allowed to differ (refer to the NVM Express Base Specification).

Figure 126: Management Endpoint - Feature Support

Feature Name ²	Feature Identifier	Support Requirements ¹	
		NVMe Storage Device	NVMe Enclosure
...			
Namespace Write Protection Config	84h	P	P

Figure 126: Management Endpoint - Feature Support

Feature Name ²	Feature Identifier	Support Requirements ¹	
		NVMe Storage Device	NVMe Enclosure
Vendor Specific	C0h to FFh	O	O
Notes: 1. O = Optional, M = Mandatory, P = Prohibited for Set Features/Optional for Get Features. 2. Refer to the NVMe Express Base Specification unless another footnote specifies otherwise. 3. Refer to the NVMe Command Set Specification. 4. Refer to the Key Value Command Set Specification.			

...