



LEGAL NOTICE:

© **Copyright 2008 to 2023 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 2023 NVM Express, Inc. ALL RIGHTS RESERVED.**” When making any such citations or references to this document you are not permitted to revise, alter, modify, make any derivatives of, or otherwise amend the referenced portion of this document in any way without the prior express written permission of NVM Express, Inc. Nothing contained in this document shall be deemed as granting you any kind of license to implement or use this document or the specification described therein, or any of its contents, either expressly or impliedly, or to any intellectual property owned or controlled by NVM Express, Inc., including, without limitation, any trademarks of NVM Express, Inc.

LEGAL DISCLAIMER:

THIS DOCUMENT AND THE INFORMATION CONTAINED HEREIN IS PROVIDED ON AN “**AS IS**” BASIS. TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, NVM EXPRESS, INC. (ALONG WITH THE CONTRIBUTORS TO THIS DOCUMENT) HEREBY DISCLAIM ALL REPRESENTATIONS, WARRANTIES AND/OR COVENANTS, EITHER EXPRESS OR IMPLIED, STATUTORY OR AT COMMON LAW, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE, VALIDITY, AND/OR NONINFRINGEMENT.

All product names, trademarks, registered trademarks, and/or servicemarks may be claimed as the property of their respective owners.

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

NVM Express® Technical Errata

Errata ID	111
Revision Date	04/17/2023
Affected Spec Ver.	NVM Express® Base Specification Revision 2.0c NVM Express® NVM Command Set Specification Revision 1.0c NVM Express® Zoned Namespace Command Set Specification Revision 1.1c NVM Express® Key Value Command Set Specification Revision 1.0c NVM Express PCIe Transport Specification Revision 1.0c NVM Express® Management Interface Specification Revision 1.2c
Corrected Spec Ver.	

Errata Author(s)

Name	Company
Judy Brock, Mike Allison, Bill Martin, Javier González, Karl Bonde Torp, Simon Lund	Samsung
Fred Knight	NetApp
Paul Suhler	Kioxia
Rory Harris	Intel

Errata Overview

This ECN updates and clarifies various text within the NVM Express Base Specification Revision 2.0b, the NVM Express NVM Command Set Specification Revision 1.0b, the NVM Express Zoned Namespace Command Set Specification Revision 1.1b, the NVM Express Key Value Command Set Specification Revision 1.0b, and the NVM Express Management Interface Specification Revision 1.2b.

Revision History

Revision Date	Change Description
4/13/2022	Initial creation
5/30/2022	Lots of updates.
5/31/2022	Reversed the bit ordering in the Read Recovery Levels Supported field. Fixed “Bit” and “Byte” in the NVM Express Key Value Command Set Specification and the NVM Express Zoned Namespace Specification.
6/1/2022	Incorporated Fred Knight (NetApp) comments
6/3/2022	Corrected the bits ordering in the AMS field and the CRMS field.
6/27/2022	Updated the CSI field in the Get Log Page command.
7/13/2022	Fixed the bit numbering in LID Supported and Effects data structure.
8/22/2022	Fixed the references in the Volatile Write Cache.
9/3/2022	Removed the renaming of a status code as is too late to change. Removed old comments. Added clarity to who is clearing. Updated the wording on the Commands

	Supported and Effects log page introduction paragraph. The CSI field in the Namespace Identification Descriptor definition was updated.
9/7/2022	Updated the definitions when CAP.CRMS field to 00b.
9/8/2022	Updates from the NVMe Technical WG meeting.
10/17/2022	Updated to new minor revisions of the specification. Added fixes to correct the name of "PCI Express Base Specification".
1/12/2023	Minor edits and removed "formatted LBA size" to another ECN. Ready for review.
3/21/2023	Integrated
4/12/2023	Integration review edits per Mike Allison and Fred Knight
4/17/2023	Editorial corrections

Description of Changes

NVM Express Base Specification 2.0c:

Backward Incompatible Changes:

- Clarified that events are not preserved across Controller Level Reset as opposed to power off conditions.
- Clarified that when the Spinup Control feature is enabled that the inhibiting is until any controller in the domain processes a Set Feature command specifying an operational state.

Editorial Changes:

- Replacing "Log Identifier" with "Log Page Identifier"
- Figure headers updated to refer to "Bits"
- Figure headers updated to refer to "Bytes"
- Updated statements that indicate a field or bit is ignore to be specific on the entity that is doing the ignoring.
- Clarified that the CC.IOCQES field is ignored by Discovery controllers as opposed to the property.
- Clarified that the CC.IOSQES field is ignored by Discovery controllers as opposed to the property.
- Updates the references of specific specification to use the correct name and capitalization.
- Updated CSTS.PP to indicate it is valid when both CC.EN and CSTS.RDY are set to '1'.
- Updated ignored bits and fields to indicate the entity doing the ignoring.
- Changed reference to "active namespace" to "attached namespace" since only NSIDs are active.
- Updated host issues protocol to use "specifies" as opposed to "indicates".
- Removed power loss notification as an Internal controller error.
- Fixed wrong figure references.
- Clarified the definition of the CSI field in the Get Log Page command to identify the Log Page Identifiers affected by that field.
- Clarified the definition of the IOS field in the LID Supported and Effects data structure for the Get Log Page command to the reference to bit 2 in the LBA field of the Identify Controller data structure.
- Clarified the Unsafe Shutdown field definition in the SMART / Health Information log page.
- Clarified that the use of the Command Set Identifier is for controllers that implement I/O Queues.
- Updated figures of bits to be in descending bit order.

- Clarified that host should use the OT bit in the Get Log Page command for the Asymmetric Namespace Access log page.
- Clarified the CSI field in the Namespace Identifier Descriptor.
- Updated the Reserved field in Command Dword 10 for the Lockdown command uses bit 7 and not bits 7:6.
- Clarified that the controller is aborting the command when the Connect Invalid Parameters status code is returned in the Connect command.
- Clarified that the DH-HMAC-CHAP_Challenge message is sent from the controller to the host.
- The CQID field in the Create I/O Submission command fixed the alignment of text for when the value of the field is within the range supported by the controller and does not specify identify an I/O Completion Queue that has been created.
- Removed black stricken text in the definition if the Volatile Write Cache field.
- Corrected spacing in the figure titled “I/O Controller – Common I/O Command Support”.
- Updated the Controller Ready Timeouts During Initialization when CAP.CRMS field is cleared 00b not not mention the CRYPTO.CRIMT field and the CRYPTO.CRWMT field.

NVM Express NVM Command Set Specification 1.0c:

Backward Incompatible Changes:

Editorial Changes:

- Replacing “Log Identifier” with “Log Page Identifier”
- Figure headers updated to refer to “Bits”
- Updated figures of bits to be in descending bit order.
- Clarified the definition of the CSI field in the Get Log Page command to identify the Log Page Identifiers affected by that field.
- Corrected the spelling of the ELBAS field in the Host Behavior Support data structure.

NVM Express Zoned Namespace Command Set Specification 1.1c:

Editorial Changes:

- Replacing “Log Identifier” with “Log Page Identifier”
- Figure headers updated to refer to “Bits”
- Removed write type status codes from the command completion section of the Compare command.
- Clarified the definition of the CSI field in the Get Log Page command to identify the Log Page Identifiers affected by that field.

NVM Express Key Value Command Set Specification 1.0c:

Editorial Changes:

- Replacing “Log Identifier” with “Log Page Identifier”
- Figure headers updated to refer to “Bits”
- Figure headers updated to refer to “Bytes”

NVM Express PCIe Transport Specification 1.0bc:

Editorial Changes:

- Correctly stating “PCI Express Base Specification” in reference to that specification.

NVM Express Management Interface Specification 1.2c:

Editorial Changes:

- Replacing “Log Identifier” with “Log Page Identifier”
- Figure headers updated to refer to “Bits”
- Figure headers updated to refer to “Bytes”

Note:

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

Description of NVM Express Base Specification 2.0c changes

<Editor: Replace all instances of “Log Identifier” with “Log Page Identifier”>

Modify a portions of section 3 as shown below:

3 NVM Express Architecture

...

3.1 NVM Controller Architecture

...

3.1.2 Controller Types

...

3.1.2.1 I/O Controller

...

3.1.2.1.1 Command Support

...

Figure 23: I/O Controller – *<add a space>*Common I/O Command Support

...

3.1.3 Controller Properties

...

3.1.3.1 Offset 0h: CAP – Controller Capabilities

...

Figure 36: Offset 0h: CAP – Controller Capabilities

Bits	Type	Reset	Description
...			

Figure 36: Offset 0h: CAP – Controller Capabilities

Bits	Type	Reset	Description								
60:59	RO	Impl Spec	Controller Ready Modes Supported (CRMS): This field indicates the ready capabilities of the controller. Refer to sections 3.5.3 and 3.5.4 for more detail.								
			<table><tr><th>Bits</th><th>Description</th></tr><tr><td>0</td><td>Controller Ready With Media Support (CRWMS): If this bit is set to '1', then the controller supports the Controller Ready With Media mode. If this bit is cleared to '0', then the controller does not support Controller Ready With Media mode. This bit shall be set to '1' on controllers compliant with versions later than NVM Express Base Specification revision 1.4.</td></tr><tr><td>1</td><td>Controller Ready Independent of Media Support (CRIMS): If this bit is set to '1', then the controller supports the Controller Ready Independent of Media mode. If this bit is cleared to '0', then the controller does not support Controller Ready Independent of Media mode.</td></tr><tr><td>0</td><td>Controller Ready With Media Support (CRWMS): If this bit is set to '1', then the controller supports the Controller Ready With Media mode. If this bit is cleared to '0', then the controller does not support Controller Ready With Media mode. This bit shall be set to '1' on controllers compliant with versions later than NVM Express Base Specification revision 1.4 NVM Express Base Specification, Revision 2.0 and later.</td></tr></table>	Bits	Description	0	Controller Ready With Media Support (CRWMS): If this bit is set to '1', then the controller supports the Controller Ready With Media mode. If this bit is cleared to '0', then the controller does not support Controller Ready With Media mode. This bit shall be set to '1' on controllers compliant with versions later than NVM Express Base Specification revision 1.4.	1	Controller Ready Independent of Media Support (CRIMS): If this bit is set to '1', then the controller supports the Controller Ready Independent of Media mode. If this bit is cleared to '0', then the controller does not support Controller Ready Independent of Media mode.	0	Controller Ready With Media Support (CRWMS): If this bit is set to '1', then the controller supports the Controller Ready With Media mode. If this bit is cleared to '0', then the controller does not support Controller Ready With Media mode. This bit shall be set to '1' on controllers compliant with versions later than NVM Express Base Specification revision 1.4 NVM Express Base Specification, Revision 2.0 and later.
			Bits	Description							
			0	Controller Ready With Media Support (CRWMS): If this bit is set to '1', then the controller supports the Controller Ready With Media mode. If this bit is cleared to '0', then the controller does not support Controller Ready With Media mode. This bit shall be set to '1' on controllers compliant with versions later than NVM Express Base Specification revision 1.4.							
1	Controller Ready Independent of Media Support (CRIMS): If this bit is set to '1', then the controller supports the Controller Ready Independent of Media mode. If this bit is cleared to '0', then the controller does not support Controller Ready Independent of Media mode.										
0	Controller Ready With Media Support (CRWMS): If this bit is set to '1', then the controller supports the Controller Ready With Media mode. If this bit is cleared to '0', then the controller does not support Controller Ready With Media mode. This bit shall be set to '1' on controllers compliant with versions later than NVM Express Base Specification revision 1.4 NVM Express Base Specification, Revision 2.0 and later.										
...											
44:37	RO	Impl Spec	Command Sets Supported (CSS): This field indicates the I/O Command Set(s) that the controller supports. ... <table><tr><th>Bits</th><th>Definition</th></tr><tr><td>...</td><td>...</td></tr></table>	Bits	Definition				
Bits	Definition										
...	...										
...											
23:19	RO	0h	Reserved								
18:17	RO	Impl Spec	Arbitration Mechanism Supported (AMS): This field is bit significant and indicates the optional arbitration mechanisms supported by the controller. If a bit is set to '1', then the corresponding arbitration mechanism is supported by the controller. Refer to section 3.4.4 for arbitration details.								
			<table><tr><th>Bits</th><th>Definition</th></tr><tr><td>0</td><td>Weighted Round Robin with Urgent Priority Class</td></tr><tr><td>1</td><td>Vendor Specific</td></tr><tr><td>0</td><td>Weighted Round Robin with Urgent Priority Class</td></tr></table>	Bits	Definition	0	Weighted Round Robin with Urgent Priority Class	1	Vendor Specific	0	Weighted Round Robin with Urgent Priority Class
			Bits	Definition							
			0	Weighted Round Robin with Urgent Priority Class							
1	Vendor Specific										
0	Weighted Round Robin with Urgent Priority Class										
....											
...											

...

3.1.3.5 Offset 14h: CC – Controller Configuration

...

Figure 46: Offset 14h: CC – Controller Configuration

Bits	Type	Reset	Description										
...													
23:20	RW/RO	0h	<p>I/O Completion Queue Entry Size (IOCQES): This field defines the I/O completion queue entry size that is used for the selected I/O Command Set(s). The required and maximum values for this field are specified in the CQES field in the Identify Controller data structure in Figure 275 for each I/O Command Set. The value is in bytes and is specified as a power of two (2^n).</p> <p>If any I/O Completion Queues exist, then write operations that change the value in this field produce undefined results.</p> <p>If the controller does not support I/O queues, then this field shall be read-only with a value of 0h.</p> <p>For Discovery controllers, this property field is reserved.</p>										
19:16	RW/RO	0h	<p>I/O Submission Queue Entry Size (IOSQES): This field defines the I/O submission queue entry size that is used for the selected I/O Command Set(s). The required and maximum values for this field are specified in the SQES field in the Identify Controller data structure in Figure 275 for each I/O Command Set. The value is in bytes and is specified as a power of two (2^n).</p> <p>If any I/O Submission Queues exist, then write operations that change the value in this field produce undefined results.</p> <p>If the controller does not support I/O queues, then this field shall be read-only with a value of 0h.</p> <p>For Discovery controllers, this property field is reserved.</p>										
...													
15:14	RW	00b	<p>Shutdown Notification (SHN): This field is used to initiate a controller shutdown when a power down condition is expected. For a normal controller shutdown, it is expected that the controller is given time to process the controller shutdown. For an abrupt shutdown, the host may not wait for the controller shutdown to complete before power is lost.</p> <p>The controller shutdown notification values are defined as:</p> <table><thead><tr><th>Value</th><th>Definition</th></tr></thead><tbody><tr><td>00b</td><td>No notification; no effect</td></tr><tr><td>01b</td><td>Normal shutdown notification</td></tr><tr><td>10b</td><td>Abrupt shutdown notification</td></tr><tr><td>11b</td><td>Reserved</td></tr></tbody></table> <p>This field should be written by host software prior to any power down condition and prior to any change of the PCI power management state. It is recommended that this field also be written prior to a warm reset (refer to the PCI Express Base Specification). To determine when the controller shutdown processing is complete, refer to CSTS.ST and CSTS.SHST. Refer to sections 3.6.1 and 3.6.2 for additional shutdown processing details.</p> <p>Other fields in this property (including the EN bit) may be modified as part of updating this field to 01b or 10b to initiate a controller shutdown. If the EN bit is cleared to '0' such that the EN bit transitions from '1' to '0', then both a Controller Reset and a controller shutdown occur.</p> <p>If an NVM Subsystem Shutdown is in progress or is being reported as completed (i.e., CSTS.ST is set to '1', and CSTS.SHST is set to 01b or 10b), then writes to this field modify the field value but have no effect. Refer to section 3.6.3 for details.</p>	Value	Definition	00b	No notification; no effect	01b	Normal shutdown notification	10b	Abrupt shutdown notification	11b	Reserved
Value	Definition												
00b	No notification; no effect												
01b	Normal shutdown notification												
10b	Abrupt shutdown notification												
11b	Reserved												
...													

3.1.3.6 Offset 1Ch: CSTS – Controller Status

...

Figure 47: Offset 1Ch: CSTS – Controller Status

Bits	Type	Reset ¹	Description
31:07	RO	0h	Reserved
06	RO	Impl Spec	Shutdown Type (ST): When CSTS.SHST is set to a non-zero value, then this bit indicates the type of shutdown reported by CSTS.SHST. If this bit is set to '1', then CSTS.SHST is reporting the state of an NVM Subsystem Shutdown. If this bit is cleared to '0', then CSTS.SHST is reporting the state of a controller shutdown. If CSTS.SHST is cleared to 00b, then this bit is should be ignored by the host.
05	RO	0b	Processing Paused (PP): This bit indicates whether the controller is processing commands. If this bit is cleared to '0', then the controller is processing commands normally. If this bit is set to '1', then the controller has temporarily stopped processing commands in order to handle an event (e.g., firmware activation). This bit is only valid when CC.EN is set to '1' and CSTS.RDY is set to '1'.
...			

...

3.1.3.10 Offset 30h: ACQ – Admin Completion Queue Base Address

...

Figure 51: Offset 30h: ACQ – Admin Completion Queue Base Address

Bits	Type	Reset	Description
...			
11:00	RO	0h	Reserved

...

3.1.3.21 Offset 68h: CRT0 – Controller Ready Timeouts

This property indicates the controller ready timeout values. This property is mandatory for controllers compliant with NVM Express Base Specification revision 2.0 and later.

Figure 62: Offset 68h: CRTO – Controller Ready Timeouts

Bits	Type	Reset	Description
31:16	RO	Impl Spec	<p>Controller Ready Independent of Media Timeout (CRIMT): If the CAP.CRMS.CRIMS bit is cleared to '0', then the controller shall clear this field to 0h and the host should ignore this field this field is not applicable and shall be cleared to 0h.</p> <p>If the CAP.CRMS.CRIMS bit is set to '1', then this field contains the worst-case time that host software should wait after CC.EN transitions from '0' to '1' for the controller to become ready and be able to successfully process all commands that do not access attached namespaces and Admin commands that do not require access to media when the controller is in Controller Ready Independent of Media mode (i.e., the CC.CRIME bit is set to '1'). Attached namespaces and media required to process Admin commands may or may not be ready within this time period (refer to section 3.5.3, section 3.5.4, and Figure 104).</p> <p>This worst-case time may be experienced after events such as an abrupt shutdown or activation of a new firmware image; typical times are expected to be much shorter. This field is in 500 millisecond units.</p> <p>The value of this field should not exceed FFh (i.e., 127.5 seconds).</p>
...			

...

3.2 NVM Subsystem Entities

...

3.2.1 Namespaces

...

3.2.1.6 NSID and Namespace Usage

...

A namespace may or may not have a relationship to a Submission Queue; this relationship is determined by the host software implementation. The controller shall support access to any attached ~~active~~ namespace from any I/O Submission Queue.

...

3.5 Controller Initialization

...

3.5.1 Memory-based Transport Controller Initialization

...

8. The host determines any I/O Command Set specific configuration information as follows:
 - a. If the CAP.CSS bit 6 is set to '1', then the host does the following:
 - i. Issue the Identify command specifying the Identify I/O Command Set data structure (CNS 1Ch); and
 - ii. Issue the Set Features command with the I/O Command Set Profile Feature Identifier (FID 19h) specifying the index of the I/O Command Set Combination (refer to Figure 289) to be enabled;

and

- b. For each I/O Command Set that is enabled (Note: the NVM Command Set is enabled if the CC.CSS field is set to 000b):
 - i. Issue the Identify command specifying the I/O Command Set specific Active Namespace ID list (CNS 07h) with the appropriate Command Set Identifier (CSI) value of that I/O Command Set; and
 - ii. For each NSID that is returned:
 1. If the enabled I/O Command Set is the NVM Command Set or an I/O Command Set based on the NVM Command Set (e.g., the Zoned Namespace Command Set) issue the Identify command specifying the Identify Namespace data structure (CNS 00h); and
 2. Issue the Identify command specifying each of the following data structures (refer to **Figure 273274**): the I/O Command Set specific Identify Namespace data structure, the I/O Command Set specific Identify Controller data structure, and the I/O Command Set independent Identify Namespace data structure;

...

3.5.4 Controller Ready Timeouts During Initialization

The CAP.CRMS field was not defined prior to NVM Express Base Specification, Revision 2.0. Controllers compliant with revisions earlier than NVM Express Base Specification, Revision 2.0 may clear the CAP.CRMS field to 00b. This section is applicable to controllers that clear the CAP.CRMS field to 00b and controllers that set CAP.CRMS to a non-zero value.

There are three controller ready timeout fields:

1. CAP.TO (refer to Figure 36);
2. CRTO.CRWMT (refer to Figure 62); and
3. CRTO.CRIMT (refer to Figure 62).

The details regarding these timeouts during controller initialization are as follows:

- a) The CAP.TO field shall be set as described in Figure 36;
- ~~b) If the CAP.CRMS field is cleared to 00b', then:~~
 - ~~i. the Controller Ready Independent of Media Timeout (CRTO.CRIMT) field is reserved;~~
 - ~~ii. the Controller Ready With Media Timeout (CRTO.CRWMT) field is reserved; and~~
 - ~~iii. the worst-case time the host should wait after the controller is enabled (i.e., CC.EN transitions from '0' to '1') for the controller to become ready (CSTS.RDY transitions from '0' to '1') is indicated by CAP.TO;~~
- b) If the CAP.CRMS field is cleared to 00b', then the worst-case time the host should wait after the controller is enabled (i.e., CC.EN transitions from '0' to '1') for the controller to become ready (CSTS.RDY transitions from '0' to '1') is indicated by CAP.TO.

...

Modify a portions of section 4 as shown below:

4 Data Structures

...

4.5 NVMe Qualified Names

...

4.5.1 Unique Identifier

The NVM Subsystem NVMe Qualified Name specified in the Identify Controller data structure (refer to Figure 275) should be used (e.g., by host software) as the unique identifier for the NVM subsystem. If the controller ~~complies with an older version of the NVM Express Base specification that does not include the NVM Subsystem NQN~~ is compliant with an NVM Express Specification prior to revision 1.2.1 (i.e., revisions where the NVM Subsystem NQN was not defined), then the PCI Vendor ID, Serial Number, and Model Number fields in the Identify Controller data structure and the NQN Starting String “nqn.2014.08.org.nvmeexpress:” may be combined by the host to form a globally unique value that identifies the NVM subsystem (e.g., for host software that uses NQNs). The method shown in Figure 138 should be used by the host to construct an NVM Subsystem NQN for older NVM subsystems that do not provide an NQN in the Identify Controller data structure. The mechanism used by the vendor to assign Serial Number and Model Number values to ensure uniqueness is outside the scope of this specification.

...

Modify a portions of section 5 as shown below:

5 Admin Command Set

...

5.2 Asynchronous Event Request command

...

If an event occurs for which reporting is enabled and there are no Asynchronous Event Request commands outstanding, the controller should retain the event information for that Asynchronous Event Type and use that information as a response to the next Asynchronous Event Request command that is received. If a Get Log Page command clears the event prior to receiving the Asynchronous Event Request command or if a ~~Controller Level Reset power-off condition~~ occurs, then a notification is not sent. If multiple events of the same type occur that have identical responses to the Asynchronous Event Request command, then those events may be reported as a single response to an Asynchronous Event Request command. If multiple events occur that are of different types or have different responses to the Asynchronous Event Request command, then the controller should retain a queue of those events for reporting in responses to subsequent Asynchronous Event Request commands.

...

5.5 Create I/O Submission Queue command

...

Figure 160: Create I/O Submission Queue – Command Dword 10

Bits	Description
31:16	Queue Size (QSIZE): This field specifies indicates the size of the Submission Queue to be created. If the size is 0h or larger than the controller supports, the controller should return an error of Invalid Queue Size. Refer to section 3.3.3.2.2. This is a 0's based value.
15:00	Queue Identifier (QID): This field specifies indicates the identifier to assign to the Submission Queue to be created. This identifier corresponds to the Submission Queue Tail Doorbell used for this command (i.e., the value y in SQyTDBL section of the NVMe over PCIe Transport Specification). This value shall not exceed the value reported in the Number of Queues feature (refer to section 5.27.1.5) for I/O Submission Queues. If the value specified is 0h, exceeds the Number of Queues reported, or corresponds to an identifier already in use, the controller should return an error of Invalid Queue Identifier.

Figure 161: Create I/O Submission Queue – Command Dword 11

Bits	Description										
31:16	<p>Completion Queue Identifier (CQID): This field specifies indicates the identifier of the I/O Completion Queue to utilize for any command completions entries associated with this Submission Queue.</p> <p>If the value specified:</p> <ul style="list-style-type: none"> a) is 0h (i.e., the Admin Completion Queue), then the controller should return an error of Invalid Queue Identifier; b) is outside the range supported by the controller, then the controller should return an error of Invalid Queue Identifier; or c) is within the range supported by the controller and does not specify identify an I/O Completion Queue that has been created, then the controller should return an error of Completion Queue Invalid. 										
15:03	Reserved										
02:01	<p>Queue Priority (QPRIO): This field specifies indicates the priority class to use for commands within this Submission Queue. This field is only used when the weighted round robin with urgent priority class is the arbitration mechanism selected. This the field shall be is ignored by the controller if weighted round robin with urgent priority class is not used. Refer to section 3.4.4.</p> <table border="1"> <thead> <tr> <th>Value</th><th>Definition</th></tr> </thead> <tbody> <tr> <td>00b</td><td>Urgent</td></tr> <tr> <td>01b</td><td>High</td></tr> <tr> <td>10b</td><td>Medium</td></tr> <tr> <td>11b</td><td>Low</td></tr> </tbody> </table>	Value	Definition	00b	Urgent	01b	High	10b	Medium	11b	Low
Value	Definition										
00b	Urgent										
01b	High										
10b	Medium										
11b	Low										
...											

...

5.9 Device Self-test command

...

Figure 171: Device Self-test Namespace Test Action

NSID Value	Description
00000000h	Specifies that the device self-test operation shall not include any namespaces, and only the controller is included as part of the device self-test operation.
00000001h to FFFFFFFEh	Specifies that the device self-test operation shall include the namespace specified by this field. If this field specifies an invalid namespace ID, then the controller shall abort the command with a status code of Invalid Namespace or Format. If this field specifies an inactive namespace ID, then the controller shall abort the command with a status code of Invalid Field in Command.
FFFFFFFFh	Specifies that the device self-test operation shall include all attached active namespaces accessible through the controller at the time the device self-test operation is started.

...

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.

Figure 188: Format NVM – Operation Scope

FNA Bit ¹	NSID	Format Operation
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 FFFFFFFFh	All namespaces that exist in the NVM subsystem.
Notes: 1. For a Format NVM command with Secure Erase, this column refers to bit 1 in the FNA field in the Identify Controller data structure (refer to Figure 275) and bit 0 in the FNA field shall be ignored by the controller . For a Format NVM command without Secure Erase, this column refers to bit 0 in the FNA field, and bit 1 in the FNA field shall be ignored by the controller . 2. If bit 3 in the FNA field is set to '1', then this value is not supported. 3. If bit 3 in the FNA field is set to '1', then this value does not occur. Refer to Figure 275.		

...

Figure 189: Format NVM – Command Dword 10

Bits	Description
...	
13:12	LBA Format Upper (LBAFU): This field specifies the most significant 2 bits of the Format Index of the User Data Format to apply to the NVM media. This corresponds to the User Data Formats indicated in the Identify command, refer to the Identify Namespace data structure and the LBA Format data structure in the applicable I/O Command Set specification. If an unsupported User Data Format is selected, the controller shall abort the command with a status code of Invalid Format. This field shall be ignored by the controller if the LBA Format Extension Enable (LBAFEE) field is cleared to 0h in the Host Behavior Support feature (refer to section 5.27.1.18). NOTE: This field applies to all User Data Formats. The original name has been retained for historical continuity.
...	

...

5.15 Get Features command

...

Figure 191: Get Features – Data Pointer

Bits	Description
127:00	Data Pointer (DPTR): This field specifies the start of the data buffer. Refer to Figure 87 for the definition of this field. If no data structure is used as part of the this specified feature (refer to the FID field), then this field shall be ignored by the controller .

...

5.16 Get Log Page command

...

Figure 201: Get Log Page – Command Dword 14

Bits	Description
31:24	Command Set Identifier (CSI): This field specifies the I/O Command Set (Refer to Figure 274) to be used by the command. This field is log page specific; refer to Figure 202 for log pages that use this field. This field shall be ignored by the controller if the specified log page does not support the use of this field (refer to Figure 202) or if the CC.CSS field is not set to 110b.
23	Offset Type (OT): If set to '1' then the Log Page Offset Lower field and the Log Page Offset Upper field specify the index into the list of data structures in the log page to be returned. If cleared to '0', then the Log Page Offset Lower field and the Log Page Offset Upper field specify the byte offset into the log page to be returned.
22:07	Reserved
06:00	UUID Index: Refer to Figure 477.

...

5.16.1 Log Specific Information

...

Figure 202: Get Log Page – Log Page Identifiers

Log Identifier	CSI ⁷	Scope	Log Page Name	Reference Section
00h	Y	Controller	Supported Log Pages	5.16.1.1
01h	N	Controller	Error Information	5.16.1.2
02h	N	Controller ¹	SMART / Health Information	5.16.1.3
		Namespace ²		
03h	N	Domain / NVM subsystem ⁶	Firmware Slot Information	5.16.1.4
04h	N	Controller	Changed Namespace List	5.16.1.5
05h	Y	Controller	Commands Supported and Effects	5.16.1.6
06h	N	Controller ³	Device Self-test ⁵	5.16.1.7
		Domain / NVM subsystem ^{4, 6}		
07h	N	Vendor Specific	Telemetry Host-Initiated ⁵	5.16.1.8
08h	N	Vendor Specific	Telemetry Controller-Initiated ⁵	5.16.1.9
09h	N	Domain / NVM subsystem ⁶	Endurance Group Information	5.16.1.10
0Ah	N	Domain / NVM subsystem ⁶	Predictable Latency Per NVM Set	5.16.1.11
0Bh	N	Domain / NVM subsystem ⁶	Predictable Latency Event Aggregate	5.16.1.12
0Ch	N	Controller	Asymmetric Namespace Access	5.16.1.13
0Dh	N	NVM subsystem	Persistent Event Log ⁵	5.16.1.14
0Eh	Refer to the NVM Command Set			
0Fh	N	Domain / NVM subsystem ⁶	Endurance Group Event Aggregate	5.16.1.15
10h	N	Domain / NVM subsystem ^{5, 6}	Media Unit Status	5.16.1.16
11h	N	Domain / NVM subsystem ⁶	Supported Capacity Configuration List	5.16.1.17

Figure 202: Get Log Page – Log Page Identifiers

Log Identifier	CSI ⁷	Scope	Log Page Name	Reference Section
12h	Y	Controller	Feature Identifiers Supported and Effects	5.16.1.18
13h	N	Controller	NVMe-MI Commands Supported and Effects	5.16.1.19
14h	Y	NVM subsystem	Command and Feature Lockdown ⁵	5.16.1.20
15h	N	Controller	Boot Partition	5.16.1.21
16h	N	Endurance Group	Rotational Media Information	5.16.1.22
17h to 6Fh	Reserved			
70h	N	Controller	Discovery	5.16.1.23
71h to 7Fh	Reserved			
80h	N	Controller	Reservation Notification	5.16.1.24
81h	N	NVM subsystem	Sanitize Status	5.16.1.25
82h to BEh	I/O Command Set Specific			
BFh	Refer to the Zoned Namespace Command Set			
C0h to FFh	Vendor specific ⁵			
<p>Key:</p> <p>Namespace = The log page contains information about a specific namespace.</p> <p>Endurance Group = The log page contains information about a specific Endurance Group.</p> <p>Controller = The log page contains information about the controller that is processing the command.</p> <p>Domain = The log page contains information about the Domain.</p> <p>NVM subsystem = The log page contains information about the NVM subsystem.</p> <p>Vendor Specific = The log page contains information that is vendor specific.</p> <p>Notes:</p> <ol style="list-style-type: none">For namespace identifiers of 0h or FFFFFFFFh.For namespace identifiers other than 0h or FFFFFFFFh.Bit 0 is cleared to '0' in the DSTO field in the Identify Controller data structure (refer to Figure 275).Bit 0 is set to '1' in the DSTO field in the Identify Controller data structure.Selection of a UUID may be supported. Refer to section 8.25.For NVM subsystems that support multiple domains (refer to the MDS bit in the Identify Controller data structure, Figure 275), Domain scope information is returned.If multiple I/O Command Sets are supported, then the CSI field is used by the log page: Y = Yes, N = No. If Yes, then refer to the definition of the log page for details on usage.				

...

5.16.1.1 Supported Log Pages (Log Identifier 00h)

...

For controllers that implement I/O Queues, the log pages that the controller supports are dependent on the I/O Command Set that is based on:

- the I/O Command Set selected in CC.CSS, if CC.CSS is not set to 110b; and
- the Command Set Identifier (CSI) field in CDW 14, if CC.CSS is set to 110b.

...

Figure 204: LID Supported and Effects Data Structure

Bits	Description
31:16	LID Specific Parameter (LIDSP): This field is specific to the log page identifier as defined in Figure 205.
15:2	Reserved

Figure 204: LID Supported and Effects Data Structure

Bits	Description
31:16	LID Specific Parameter (LIDSP): This field is specific to the log page identifier as defined in Figure 205.
1	Index Offset Supported (IOS): If this bit is set to '1', then the controller supports an index offset for this LID in a Get Log Page command (i.e., the OT bit in the Get Log Page command is allowed to be set to '1'). If this bit is cleared to '0', then the controller does not support an index offset for this LID in a Get Log Page command (i.e., the OT bit in the Get Log Page command is only allowed to be cleared to '0'). If the controller does not support extended data for the Get Log Page command (refer to bit 2 in the LPA field in Figure 275), then this bit shall be cleared to '0'.
0	LID Supported (LSUPP): If this bit is set to '1', then the controller supports this LID for a Get Log Page command. If this bit is cleared to '0', then the controller does not support this LID for a Get Log Page command. Refer to section 3.1.2 for the LID support requirements for each controller type.

...

5.16.1.3 SMART / Health Information (Log Identifier 02h)

...

Figure 207: SMART / Health Information Log Page

Bytes	Description
...	
159:144	Unsafe Shutdowns: Contains the number of unsafe shutdowns. This count is incremented when the controller does not report it is safe to power down prior to loss of main power. If the Controller Power Scope (i.e., CAP.CPS) field is cleared to 00b (i.e., Not Reported) or set to 01b (i.e., Controller scope), then the controller reports that it is safe to power down the controller when a controller shutdown processing is complete (i.e., CSTS.ST bit is cleared to '0' and CSTS.SHST field is set to 10b). If CAP.CPS field is set to 10b (i.e., Domain scope), then the controller reports that it is safe to power down the domain when NVM Subsystem Shutdown processing is complete (i.e., CSTS.ST bit is set to '1' and CSTS.SHST field is set to 10b). If CAP.CPS field is set to 11b (i.e., NVM subsystem scope), then the controller reports that it is safe to power down the NVM subsystem when NVM Subsystem Shutdown processing is complete (i.e., CSTS.ST bit is set to '1' and CSTS.SHST field is set to 10b).
...	

...

5.16.1.6 Commands Supported and Effects (Log Identifier 05h)

This log page is used to describe the commands that the controller supports and the effects of those commands on the state of the NVM subsystem. The log page is 4,096 bytes in size. There is one Commands Supported and Effects data structure per Admin command opcode value. For controllers that implement I/O Queues, there is ~~and~~ one Commands Supported and Effects data structure per I/O command opcode value for a specified I/O Command Set based on:

- the I/O Command Set selected in CC.CSS, if CC.CSS is not set to 110b; and
- the Command Set Identifier field in CDW 14, if CC.CSS is set to 110b.

...

5.16.1.7 Device Self-test (Log Identifier 06h)

This log page is used to indicate:

- a) the status of any device self-test operation in progress and the percentage complete of that operation; and
- b) the results of the last 20 device self-test operations.

The Self-test Result Data Structure contained in the Newest Self-test Result Data Structure field is always the result of the last completed or aborted self-test operation. The next Self-test Result Data Structure field in the Device Self-test log page contains the results of the second newest self-test operation and so on. If fewer than 20 self-test operations have completed or been aborted, then the Device Self-test Status field shall be set to Fh in the unused Self-test Result Data Structure fields and all other fields in that Self-test Result Data Structure ~~should be~~ **are** ignored **by the host**.

Figure 212: Device Self-test Log Page

Bytes	Description														
00	<p>Current Device Self-Test Operation: This field defines the current device self-test operation. Bits 7:4 are reserved.</p> <p>Bits 3:0 indicates the status of the current device self-test operation as defined in the following table. If a device self-test operation is in process (i.e., this field is set to 1h or 2h), then the controller shall not set this field to 0h until a new Self-test Result Data Structure is created (i.e., if a device self-test operation completes or is aborted, then the controller shall create a Self-test Result Data Structure prior to setting this field to 0h).</p> <table border="1"> <thead> <tr> <th>Value</th><th>Definition</th></tr> </thead> <tbody> <tr> <td>0h</td><td>No device self-test operation in progress</td></tr> <tr> <td>1h</td><td>Short device self-test operation in progress</td></tr> <tr> <td>2h</td><td>Extended device self-test operation in progress</td></tr> <tr> <td>3h to Dh</td><td>Reserved</td></tr> <tr> <td>Eh</td><td>Vendor specific</td></tr> <tr> <td>Fh</td><td>Reserved</td></tr> </tbody> </table>	Value	Definition	0h	No device self-test operation in progress	1h	Short device self-test operation in progress	2h	Extended device self-test operation in progress	3h to Dh	Reserved	Eh	Vendor specific	Fh	Reserved
Value	Definition														
0h	No device self-test operation in progress														
1h	Short device self-test operation in progress														
2h	Extended device self-test operation in progress														
3h to Dh	Reserved														
Eh	Vendor specific														
Fh	Reserved														
01	<p>Current Device Self-Test Completion: This field defines the completion status of the current device self-test. Bit 7 is reserved.</p> <p>Bits 6:0 indicates the percentage of the device self-test operation that is complete (e.g., a value of 25 indicates that 25% of the device self-test operation is complete and 75% remains to be tested). If bits 3:0 in the Current Device Self-Test Operation field are cleared to 0h (indicating there is no device self-test operation in progress), then this field is should be ignored by the host.</p>														
03:02	Reserved														

...

Figure 213: Self-test Result Data Structure

Bytes	Description
...	
01	<p>Segment Number: This field indicates the segment number (refer to section 8.6) where the first self-test failure occurred. If Device Self-test Status field bits [3:0] are not set to 7h, then this field should be ignored by the host.</p>
...	

...

5.16.1.11 Predictable Latency Per NVM Set (Log Identifier 0Ah)

...

Figure 218: Get Log Page – Predictable Latency Per NVM Set Log

Bytes	Description																								
...																									
01	Reserved																								
03:02	<p>Event Type: This field specifies the event(s) that occurred for the NVM Set indicated. Multiple bits may be set to '1'. All bits are cleared to '0' after the log page is read with Retain Asynchronous Event bit cleared to '0'.</p> <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> <tr> <td>16</td><td>Autonomous transition from DTWIN to NDWIN due to typical or maximum value exceeded</td></tr> <tr> <td>17:03</td><td>Reserved</td></tr> <tr> <td>02</td><td>DTWIN Time Warning</td></tr> <tr> <td>01</td><td>DTWIN Writes Warning</td></tr> <tr> <td>00</td><td>DTWIN Reads Warning</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	16	Autonomous transition from DTWIN to NDWIN due to typical or maximum value exceeded	17:03	Reserved	02	DTWIN Time Warning	01	DTWIN Writes Warning	00	DTWIN Reads Warning
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																								
16	Autonomous transition from DTWIN to NDWIN due to typical or maximum value exceeded																								
17:03	Reserved																								
02	DTWIN Time Warning																								
01	DTWIN Writes Warning																								
00	DTWIN Reads Warning																								
...																									

...

5.16.1.13 Asymmetric Namespace Access (Log Identifier 0Ch)

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

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

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

If the Index Offset Supported bit is set to '1' in the LID Supported and Effects data structure for this log page (refer to Figure 204), then: the entry data structure that is indexed is an ANA Group Descriptor (e.g., specifying an index offset of 2 returns this log page starting at the offset of ANA Group Descriptor 1). [Due to the complexity of using a byte offset with a variable sized ANA Group Descriptor, the host should use an index offset by setting the OT bit to '1' in Command Dword 14 of the Get Log Page command.](#)

...

5.16.1.14 Persistent Event Log (Log Identifier 0Dh)

...

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
...	
05h	<p>PCIe Link Not Active: Indicates that the Data Link Control and Management State Machine (refer to PCI Express Base Specification) has transitioned out of the DL_Active state without a corresponding event (e.g., without an indication from the host that the link is to be disabled).</p> <p>This NVM subsystem hardware error event does not contain additional hardware error information.</p>
...	

...

Figure 234: Additional Hardware Error Information for correctable and uncorrectable PCIe errors

Bytes	Value
01:00	PCIe Device Status Register: Contains the contents of the PCI Device Status Register (refer to the PCI Express Base Specification) at the time of the event.
02	<p>Bits 7:1 Reserved</p> <p>Bit 0 PCIe AER Supported: set to '1' indicates that PCIe AER (refer to the PCI Express Base Specification) is supported and that the PCIe AER Error Status field, PCIe AER Error Mask field, PCIe AER Header Log Register field, and the PCIe AER TLP Prefix Log Register field are reported. Bit 0 cleared to '0' indicates that PCIe AER is not supported and that the PCIe AER Error Status field, PCIe AER Error Mask field, PCIe AER Header Log Register field, and PCIe AER TLP Prefix Log Register field are not reported (i.e., bytes 79:16 are not reported).</p>
...	

...

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).

For controllers that implement I/O Queues, the ~~the~~ features that the controller supports are dependent on the I/O Command Set that is based on:

- the I/O Command Set selected in CC.CSS, if CC.CSS is not set to 110b; and
- the Command Set Identifier (CSI) field in CDW 14, if CC.CSS is set to 110b.

...

5.16.1.20 Command and Feature Lockdown (Log Identifier 14h)

...

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

...

5.17 Identify command

...

5.17.2 Identify Data Structures

...

5.17.2.1 Identify Controller Data Structure (CNS 01h)

...

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

Bytes	I/O ¹	Admin ¹	Disc ¹	Description
...				

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

Bytes	I/O ¹	Admin ¹	Disc ¹	Description																																																															
101:100	O	O	R	Read Recovery Levels Supported (RRLS): If Read Recovery Levels (RRL) are supported, then this field shall be supported. If a bit is set to '1', then the corresponding Read Recovery Level is supported. If a bit is cleared to '0', then the corresponding Read Recovery Level is not supported.																																																															
				Bits	Definition	0	Read Recovery Level 0	1	Read Recovery Level 1	2	Read Recovery Level 2	3	Read Recovery Level 3	4	Read Recovery Level 4 – Default ¹	5	Read Recovery Level 5	6	Read Recovery Level 6	7	Read Recovery Level 7	8	Read Recovery Level 8	9	Read Recovery Level 9	10	Read Recovery Level 10	11	Read Recovery Level 11	12	Read Recovery Level 12	13	Read Recovery Level 13	14	Read Recovery Level 14	15	Read Recovery Level 15 – Fast Fail ¹	14	Read Recovery Level 14	13	Read Recovery Level 13	12	Read Recovery Level 12	11	Read Recovery Level 11	10	Read Recovery Level 10	9	Read Recovery Level 9	8	Read Recovery Level 8	7	Read Recovery Level 7	6	Read Recovery Level 6	5	Read Recovery Level 5	4	Read Recovery Level 4 – Default ¹	3	Read Recovery Level 3	2	Read Recovery Level 2	1	Read Recovery Level 1	0	Read Recovery Level 0
				Bits	Definition																																																														
				0	Read Recovery Level 0																																																														
				1	Read Recovery Level 1																																																														
				2	Read Recovery Level 2																																																														
				3	Read Recovery Level 3																																																														
				4	Read Recovery Level 4 – Default ¹																																																														
				5	Read Recovery Level 5																																																														
				6	Read Recovery Level 6																																																														
				7	Read Recovery Level 7																																																														
				8	Read Recovery Level 8																																																														
				9	Read Recovery Level 9																																																														
				10	Read Recovery Level 10																																																														
				11	Read Recovery Level 11																																																														
				12	Read Recovery Level 12																																																														
				13	Read Recovery Level 13																																																														
				14	Read Recovery Level 14																																																														
				15	Read Recovery Level 15 – Fast Fail ¹																																																														
				14	Read Recovery Level 14																																																														
				13	Read Recovery Level 13																																																														
				12	Read Recovery Level 12																																																														
				11	Read Recovery Level 11																																																														
				10	Read Recovery Level 10																																																														
				9	Read Recovery Level 9																																																														
				8	Read Recovery Level 8																																																														
				7	Read Recovery Level 7																																																														
				6	Read Recovery Level 6																																																														
				5	Read Recovery Level 5																																																														
				4	Read Recovery Level 4 – Default ¹																																																														
				3	Read Recovery Level 3																																																														
				2	Read Recovery Level 2																																																														
				1	Read Recovery Level 1																																																														
0	Read Recovery Level 0																																																																		
NOTE: 1. If Read Recovery Levels are supported, then this bit shall be set to '1'.																																																																			

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

Bytes	I/O ¹	Admin ¹	Disc ¹	Description																		
315:312	O	O	R	Replay Protected Memory Block Support (RPMBs): This field indicates if the controller supports one or more Replay Protected Memory Blocks (RPMBs) and the capabilities. Refer to section 8.18.																		
				<table><tr><th>Bits</th><th>Description</th></tr><tr><td>31:24</td><td>Access Size: If the Number of RPMB Units field is non-zero, then this field indicates the maximum number of 512B units of data that may be read or written per RPMB access by Security Send or Security Receive commands for the controller. This is a 0's based value. A value of 0h indicates support for one unit of 512B of data. If the Number of RPMB Units field is 0h, then this field shall be ignored by the host.</td></tr><tr><td>23:16</td><td>Total Size: If the Number of RPMB Units field is non-zero, then this field indicates the number of 128 KiB units of data in each RPMB supported in the controller. This is a 0's based value. A value of 0h indicates support for one unit of 128 KiB of data. If the Number of RPMB Units field is 0h, this field shall be ignored by the host.</td></tr><tr><td>15:06</td><td>Reserved</td></tr><tr><td>05:03</td><td>Authentication Method: This field indicates the authentication method used to access all RPMBs in the controller. The values for this field are:<table><tr><th>Value</th><th>Definition</th></tr><tr><td>000b</td><td>HMAC SHA-256 (refer to RFC 6234)</td></tr><tr><td>001b to 111b</td><td>Reserved</td></tr></table></td></tr><tr><td>02:00</td><td>Number of RPMB Units: This field indicates the number of RPMB targets the controller supports. All RPMB targets supported shall have the same capabilities as defined in the RPMBs field. A value of 0h indicates the controller does not support Replay Protected Memory Blocks. If this value is non-zero, then the controller shall support the Security Send and Security Receive commands.</td></tr></table>	Bits	Description	31:24	Access Size: If the Number of RPMB Units field is non-zero, then this field indicates the maximum number of 512B units of data that may be read or written per RPMB access by Security Send or Security Receive commands for the controller. This is a 0's based value. A value of 0h indicates support for one unit of 512B of data. If the Number of RPMB Units field is 0h, then this field shall be ignored by the host .	23:16	Total Size: If the Number of RPMB Units field is non-zero, then this field indicates the number of 128 KiB units of data in each RPMB supported in the controller. This is a 0's based value. A value of 0h indicates support for one unit of 128 KiB of data. If the Number of RPMB Units field is 0h, this field shall be ignored by the host .	15:06	Reserved	05:03	Authentication Method: This field indicates the authentication method used to access all RPMBs in the controller. The values for this field are: <table><tr><th>Value</th><th>Definition</th></tr><tr><td>000b</td><td>HMAC SHA-256 (refer to RFC 6234)</td></tr><tr><td>001b to 111b</td><td>Reserved</td></tr></table>	Value	Definition	000b	HMAC SHA-256 (refer to RFC 6234)	001b to 111b	Reserved	02:00	Number of RPMB Units: This field indicates the number of RPMB targets the controller supports. All RPMB targets supported shall have the same capabilities as defined in the RPMBs field. A value of 0h indicates the controller does not support Replay Protected Memory Blocks. If this value is non-zero, then the controller shall support the Security Send and Security Receive commands.
				Bits	Description																	
				31:24	Access Size: If the Number of RPMB Units field is non-zero, then this field indicates the maximum number of 512B units of data that may be read or written per RPMB access by Security Send or Security Receive commands for the controller. This is a 0's based value. A value of 0h indicates support for one unit of 512B of data. If the Number of RPMB Units field is 0h, then this field shall be ignored by the host .																	
				23:16	Total Size: If the Number of RPMB Units field is non-zero, then this field indicates the number of 128 KiB units of data in each RPMB supported in the controller. This is a 0's based value. A value of 0h indicates support for one unit of 128 KiB of data. If the Number of RPMB Units field is 0h, this field shall be ignored by the host .																	
				15:06	Reserved																	
05:03	Authentication Method: This field indicates the authentication method used to access all RPMBs in the controller. The values for this field are: <table><tr><th>Value</th><th>Definition</th></tr><tr><td>000b</td><td>HMAC SHA-256 (refer to RFC 6234)</td></tr><tr><td>001b to 111b</td><td>Reserved</td></tr></table>	Value	Definition	000b	HMAC SHA-256 (refer to RFC 6234)	001b to 111b	Reserved															
Value	Definition																					
000b	HMAC SHA-256 (refer to RFC 6234)																					
001b to 111b	Reserved																					
02:00	Number of RPMB Units: This field indicates the number of RPMB targets the controller supports. All RPMB targets supported shall have the same capabilities as defined in the RPMBs field. A value of 0h indicates the controller does not support Replay Protected Memory Blocks. If this value is non-zero, then the controller shall support the Security Send and Security Receive commands.																					
...																						
525	M	M	R	Volatile Write Cache (VWC): This field indicates attributes related to the presence of a volatile write cache in the controller. Bits 7:3 are reserved. Bits 2:1 indicate Flush command behavior (refer to the Flush Command section 7.1 of the NVM Command Set Specification 7.1) if the NSID value is set to FFFFFFFFh as follows: ... If a volatile write cache is present, then the host controls whether the volatile write cache is enabled with a Set Features command specifying the Volatile Write Cache feature identifier (refer to section 5.27.1.4). The Flush command (refer to the Flush Command section 7.1 of the NVM Command Set Specification 7.1) is used to request that the contents of a volatile write cache be made non-volatile.																		
...																						

...

15.17.2.3 Namespace Identification Descriptor list (CNS 03h)

...

Figure 277: Identify – Namespace Identification Descriptor

Bytes	Description		
00	Namespace Identifier Type (NIDT): This field indicates the data type contained in the Namespace Identifier field and the length of that type as defined in the following table.		
	Value	Length (NIDL)	Definition
	...		
	4h	1h	Command Set Identifier (CSI): The NID field contains the Command Set Identifier that indicates the I/O Command Set that is associated with operates on this namespace. Refer to Figure 274.
	5h to FFh		Reserved
...			

...

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

...

Figure 290: I/O Command Set Vector

Bits	Description
...	

...

5.19 Lockdown command

...

<Editor: Merge the figure title into the Table>

Figure 291: Lockdown – Command Dword 10

Bits	Description										
31:16	Reserved										
15:08	Opcode or Feature Identifier (OFI): This field specifies the command opcode or Set Features Feature Identifier identified by the Scope field.										
07:06	Reserved										
06:05	Interface (IFC): This field identifies the interfaces affected by this command. The actions of this command apply if a command is received on the specified interfaces.										
	<table><tr><th>Value</th><th>Affected Interfaces</th></tr><tr><td>00b</td><td>Admin Submission Queue</td></tr><tr><td>01b</td><td>Admin Submission Queue and out-of-band on a Management Endpoint</td></tr><tr><td>10b</td><td>Out-of-band on a Management Endpoint</td></tr><tr><td>11b</td><td>Reserved</td></tr></table>	Value	Affected Interfaces	00b	Admin Submission Queue	01b	Admin Submission Queue and out-of-band on a Management Endpoint	10b	Out-of-band on a Management Endpoint	11b	Reserved
	Value	Affected Interfaces									
	00b	Admin Submission Queue									
	01b	Admin Submission Queue and out-of-band on a Management Endpoint									
10b	Out-of-band on a Management Endpoint										
11b	Reserved										
...											

...

5.27 Set Features command

...

5.27.1 Feature Specific Information

...

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 Log Page 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.																							
	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.	16	Autonomous transition from DTWIN to NDWIN due to typical or maximum value exceeded.	17 to 31	Reserved	32	DTWIN Time Warning	33	DTWIN Writes Warning	34	DTWIN Reads Warning
	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.																						
	16	Autonomous transition from DTWIN to NDWIN due to typical or maximum value exceeded.																						
	17 to 31	Reserved																						
	32	DTWIN Time Warning																						
	33	DTWIN Writes Warning																						
	34	DTWIN Reads Warning																						
...																								

...

5.27.1.21 I/O Command Set Profile (Feature Identifier 19h)

...

Figure 354: I/O Command Set Profile – Command Dword 11

...

5.27.1.23 Host Metadata (Feature Identifier 7Dh), (Feature Identifier 7Eh), (Feature Identifier 7Fh)

...

Figure 361: Metadata Element Descriptor

Bits	Description
...	

...

Modify a portions of section 6 as shown below:

...

6 Fabrics Command Set

...

6.1 Authentication Receive Command and Response

...

Figure 376: Authentication Receive Command – Submission Queue Entry

...

6.3 Connect Command and Response

...

If an NVM subsystem supports DH-HMAC-CHAP authentication (refer to section 8.13.6), then the Host NQN and the NVM Subsystem NQN parameters in a Connect command are required to be different. If the Host NQN and the NVM Subsystem NQN parameters in a Connect command are identical and the NVM subsystem supports DH-HMAC-CHAP authentication, then the controller shall abort the command with a status code of Connect Invalid Host.

...

Modify a portions of section 8 as shown below:

8 Extended Capabilities

...

8.13 NVMe over Fabrics Secure Channel and In-band Authentication

...

8.13.5 DH-HMAC-CHAP Protocol

...

8.13.5.3 DH-HMAC-CHAP_Challenge Message

The DH-HMAC-CHAP_Challenge message is sent from ~~for~~ the controller to the host. The format of the DH-HMAC-CHAP_Challenge message is shown in Figure 448.

...

8.15 Power Management

...

8.15.4 Runtime D3 Transitions

...

In this specification, RTD3 refers to the D3_{cold} power state described in the PCI Express Base Specification. RTD3 does not include the PCI Express D3_{hot} power state because main power is not removed from the controller in the D3_{hot} power state. Refer to the PCI Express Base Specification for details on the D3_{hot} power state and the D3_{cold} power state.

...

8.20 Rotational Media

...

If:

- a) a domain contains an Endurance Group that stores data on rotational media;
- b) that domain processes an NVM Subsystem Reset; and
- c) the Spinup Control feature (refer to section 5.27.1.22) is:
 - a. disabled, then initial spinup for all such Endurance Groups in that domain shall be initiated; and
 - b. enabled, then initial spinup for all such Endurance Groups in that domain shall be inhibited during processing of the NVM Subsystem Reset until any the controller within that domain processes a Set Features (Power Management) command that specifies an operational power state.

...

Modify a portion of section 9 as shown below:

9 Error Reporting and Recovery

...

9.4 Internal Controller Error Handling

Errors such as a DRAM failure ~~or power loss notification~~ indicate that a controller level failure has occurred during the processing of a command. The status code of the completion queue entry should indicate an Internal Error status code. Host software shall ignore any data transfer associated with the command. The host may choose to re-submit the command or indicate an error to the higher level software.

...

Description of NVM Express NVM Command Set specification 1.0c changes

<Editor: Replace all instances of “Log Identifier” with “Log Page Identifier”>

Modify a portions of section 2 as shown below:

2 NVM Command Set Model

...

2.1 Theory of operation

...

2.1.5 End-to-end Protection Information

...

Figure 9: Protection Information Field Definition

Bits	Description				
...					
02:00	Protection Information Check (PRCHK): The protection information check field specifies the fields that shall be checked as part of end-to-end data protection processing. If the namespace is not formatted to use end-to-end protection information, then this field shall be ignored by the controller. Refer to section 5.2. <table><tr><th>Bits</th><th>Definition</th></tr><tr><td>...</td><td></td></tr></table>	Bits	Definition	...	
Bits	Definition				
...					

...

Modify a portions of section 3 as shown below:

3 I/O Commands for the NVM Command Set

...

3.2 NVM Command Set Commands

...

3.2.3 Dataset Management command

...

Figure 40: Dataset Management – Command Dword 11

Bits	Description
...	

...

3.2.7 Write Uncorrectable command

...

3.2.7.1 Command Completion

...

Figure 70: Write Uncorrectable – Command Specific Status Values

Bit Value	Description
...	

...

Modify a portions of section 4 as shown below:

4 Admin Commands for the NVM Command Set

...

4.1 Admin Command behavior for the NVM Command Set

...

4.1.3 Get Features & Set Features commands

...

4.1.3.4 Host Behavior Support (Feature Identifier 16h)

...

Figure 86: Host Behavior Support – Data Structure

Bytes	Description
02	LBA Format Extension Enable (LBAFEE): This field allows the host to specify support for the extended LBA formats (refer to the EBLAS ELBAS field in the Identify Controller data structure in the NVM Express Base Specification). If this field is set to 1h and the ELBAS field is set to '1', then the controller: ...

...

4.1.4 Get Log Page command

...

Figure 89: Get Log Page – Log Page Identifiers

Log Identifier	CSI ¹	Scope	Log Page Name	Reference
01h	N	Controller	Error Information	4.1.4.1
02h	N	Controller or NVM Subsystem	SMART / Health Information	4.1.4.2

Figure 89: Get Log Page – Log Page Identifiers

Log Identifier	CSI ¹	Scope	Log Page Name	Reference
06h	N	Controller	Device Self-test	4.1.4.3
0Eh	N	Controller	LBA Status Information	4.1.4.4
Note: 1. If multiple I/O Command Sets are supported (refer to the NVM Express Base Specification), then the CSI field is used by the log page: Y = Yes, N = No. If Yes, then refer to the definition of the log page for details on usage.				

...

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
07:00	M	Namespace Size (NSZE): This field indicates the total size of the namespace in logical blocks. A namespace of size n consists of LBA 0 through LBA $(n - 1)$. The number of logical blocks is based on the formatted LBA size.
...		

...

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

<Editor: Replace all instances of “Log Identifier” with “Log Page Identifier”>

Modify a portions of section 3 as shown below:

3 I/O Commands for the Zoned Namespace Command Set

...

3.3 NVM Command Set I/O Commands

...

3.3.1 Compare command

...

3.3.1.1 Command Completion

...

Figure 13: Compare – Command Specific Status Values

Value	Description
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.
BCh	Zone Invalid Write: The write to a zone was not at the write pointer.
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 portions of section 4 as shown below:

4 Admin Commands for the Zoned Namespace Command Set

...

4.1 Admin Command behavior for the Zoned Namespace Command Set

...

4.1.3 Get Features and Set Features Commands

...

Figure 44: Asynchronous Event Configuration – Command Dword 11

Bits	Description
...	

...

4.1.4 Get Log Page Command

...

Figure 45: Log Page Identifiers

Log Identifier	CSI ²	Scope	Log Page Name	Reference
BFh	Y	Namespace ¹	Changed Zone List	4.1.4.1
KEY: Namespace = The log page contains information about a specific zoned namespace.				
NOTES: 1. For namespace identifiers other than 0h or FFFFFFFFh. 2. If multiple I/O Command Sets are supported (refer to the NVM Express Base Specification), then the CSI field is used by the log page: Y = Yes, N = No. If Yes, refer to the definition of the log page for details on usage.				

...

Description of NVM Express Key Value Command Set specification 1.1c changes

<Editor: Please all instances of “Log Identifier” with “Log Page Identifier”>

Modify a portions of section 3 as shown below:

3 I/O Commands for the Key Value Command Set

...

3.2 Key Value Command Set Commands

...

3.2.1 Delete command

...

Figure 6: Delete – Command Dword 11

Bits	Description
...	

Figure 7: Delete – Command Dword 2 and Command Dword 3

Bits	Description
...	

Figure 8: Delete – Command Dword 14 and Command Dword 15

Bits	Description
...	

...

3.2.2 List command

..

Figure 10: List – Command Dword 10

Bits	Description
...	

Figure 11: List – Command Dword 11

Bits	Description
...	

Figure 12: List – Command Dword 2 and Command Dword 3

Bits	Description
...	

Figure 13: List – Command Dword 14 and Command Dword 15

Bits	Description
...	

...

3.2.2.2 List command return data structure

...

Figure 15: List – Return data structure

Bytes	Description
...	

Figure 16: Key data structure

Bytes	Description
...	

...

3.2.3 Retrieve command

...

Figure 17: Retrieve – Data Pointer

Bits	Description
...	

Figure 18: Retrieve – Command Dword 10

Bits	Description
...	

Figure 19: Retrieve – Command Dword 11

Bits	Description
...	

Figure 20: Retrieve – Command Dword 2 and Command Dword 3

Bits	Description
...	

Figure 21: Retrieve –Command Dword 14 and Command Dword 15

Bits	Description
...	

...

3.2.4 Exist command

...

Figure 23: Exist – Command Dword 11

Bits	Description
...	

Figure 24: Exist – Command Dword 2 and Command Dword 3

Bits	Description
...	

Figure 25: Exist –<delete extra space> Command Dword 14 and Command Dword 15

Bits	Description
...	

...

3.2.5 Store command

...

Figure 27: Store – Data Pointer

Bits	Description
...	

Figure 28: Store – Command Dword 10

Bits	Description
...	

Figure 29: Store – Command Dword 11

Bits	Description
...	

Figure 30: Store – Command Dword 2 and Command Dword 3

Bits	Description
...	

Figure 31: Store –Command Dword 14 and Command Dword 15

Bits	Description
...	

...

Description of NVM Express PCIe Transport Specification 1.0c changes

Modify a portions of section 3 as shown below:

3 Transport Binding

...

3.8 Transport Specific Content

3.8.1 PCI Express Type 0/1 Common Configuration Space

Figure 10 summarizes the organization of the Type 0/1 Common Configuration Space defined in the PCI Express Base Specification. The reference information in this section does not contain all PCI Express requirements. Refer to the PCI Express Base Specification for more information.

...

Description of NVM Express Management Interface Specification 1.2c changes

Modify a portions of section 4 as shown below:

4 Message Servicing Model

...

4.1 NVMe-MI Messages

...

4.1.2 Response Messages

...

4.1.2.3 More Processing Required Response

...

Figure 33: More Processing Required Response Fields

Bytes	Description
...	

...

Modify a portions of section 5 as shown below:

5 Management Interface Command Set

...

5.11 Shutdown

...

Figure 109: Shutdown - NVMe Management Dword 0

Bits	Description
...	

...

Modify a portions of section 8 as shown below:

8 Management Architecture

...

8.2 Vital Product Data

...

Figure 148: VPD Elements

Bytes	Name
...	

...

8.2.5 Topology MultiRecord Area

...

8.2.5.8 FRU Information Device Element Descriptor

...

Figure 174: FRU Information Device Element Descriptor

Byte Offset	Factory Default	Description				
...						
03	A6h/A7h or 0h for NVM Storage Devices A4h/A5h or 0h for Carriers	SMBus/I2C Address Info: If the NVMe Storage Device contains an SMBus/I2C port, then this field indicates the default SMBus/I2C addressing per the table below; else, this field shall be cleared to 0h. <table><tr><th>Bits</th><th>Description</th></tr><tr><td colspan="2">...</td></tr></table>	Bits	Description	...	
Bits	Description					
...						
04	Impl Spec	SMBus/I2C Capabilities: If the NVM Storage Device contains an SMBus/I2C port, then this field indicates the SMBus/I2C capabilities per the table below; else, this field shall be cleared to 0h. <table><tr><th>Bits</th><th>Description</th></tr><tr><td colspan="2">...</td></tr></table>	Bits	Description	...	
Bits	Description					
...						
...						

...

Modify a portions of appendix A as shown below:

Appendix A Technical Note: NVM Express Basic Management Command

...

Figure 176: Subsystem Management Data Structure

Command Code	Offset (byte)	Description				
0	...					
	06	Current Power (Optional): This field reports the current NVM Subsystem power consumption. If both bit mapped fields are cleared to 0h, then this field is not reported.				
		<table><tr><th>Bits</th><th>Definition</th></tr><tr><td>...</td><td></td></tr></table>	Bits	Definition	...	
		Bits	Definition			
...						

Figure 176: Subsystem Management Data Structure

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

...