



**LEGAL NOTICE:**

© Copyright 2007 to 2021 NVM Express™, Inc. ALL RIGHTS RESERVED.

This NVM Express revision 1.4 technical proposal is proprietary to the NVM Express, Inc. (also referred to as "Company") and/or its successors and assigns.

**NOTICE TO USERS WHO ARE NVM EXPRESS, INC. MEMBERS:** Members of NVM Express, Inc. have the right to use and implement this NVM Express revision 1.4 technical proposal subject, however, to the Member's continued compliance with the Company's Intellectual Property Policy and Bylaws and the Member's Participation Agreement.

**NOTICE TO NON-MEMBERS OF NVM EXPRESS, INC.:** If you are not a Member of NVM Express, Inc. and you have obtained a copy of this document, you only have a right to review this document or make reference to or cite this document. Any such references or citations to this document must acknowledge NVM Express, Inc. copyright ownership of this document. The proper copyright citation or reference is as follows: "© 2007 to 2021 NVM Express, Inc. ALL RIGHTS RESERVED." When making any such citations or references to this document you are not permitted to revise, alter, modify, make any derivatives of, or otherwise amend the referenced portion of this document in any way without the prior express written permission of NVM Express, Inc. Nothing contained in this document shall be deemed as granting you any kind of license to implement or use this document or the specification described therein, or any of its contents, either expressly or impliedly, or to any intellectual property owned or controlled by NVM Express, Inc., including, without limitation, any trademarks of NVM Express, Inc.

**LEGAL DISCLAIMER:**

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

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

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

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

## NVM Express Technical Proposal for New Feature

Technical Proposal ID	4056d
Change Date	2021-02-04
Builds on Specification(s)	NVM Express specification 1.4 NVM Express Management Interface Specification 1.1

### Technical Proposal Author(s)

Name	Company
Peter Onufryk	Intel
Christoph Hellwig	Western Digital
Bill Martin	Samsung

This technical proposal adds support for namespace types. NVMe currently only supports namespaces that contain traditional logical blocks that may be randomly read and written. In some applications there is value in exposing namespaces that contain logical blocks that have special access rules (e.g., must be written sequentially), namespaces that contain logical blocks that are used to perform data transformations, or namespaces that contain things other than logical blocks. The purpose of this technical proposal is to extend the NVMe architecture to support different types of namespaces. This extension is fully backward compatible with the existing NVMe specification.

TP 4056d replaced TP 4056c that was never published.

### Revision History

Revision Date	Change Description
2019-01-16	<ul style="list-style-type: none"><li>Initial version.</li></ul>
2019-05-15	<ul style="list-style-type: none"><li>New version with significant changes from initial version based on extensive workgroup discussion and feedback.</li></ul>
2019-05-23	<ul style="list-style-type: none"><li>Updated based on workgroup member feedback.</li><li>Added Command Set Restriction feature.</li><li>Added reserved space in conventional logical block Identify Namespace data structure for derived namespaces.</li></ul>
2019-05-29	<ul style="list-style-type: none"><li>Changed name of command set restriction feature to command set profile and updated behavior to align with a profile rather than a restriction.</li><li>Changed name of Command Set Restricted error to Command Set Not Supported in Command Set Profile.</li><li>Eliminated Command Set Combination Not Supported error since this error is no longer possible given the new command set profile feature behavior.</li></ul>

2019-06-05	<ul style="list-style-type: none"> <li>Added Command Set Combination Rejected error to have an error status that clearly distinguishes the case of an invalid command set combination verses trying to use a namespace combination that is incompatible with a currently attached namespace.</li> <li>Removed reserved space in Conventional Logical Block Namespace – Identify Namespace Data Structure for derived namespaces based on workgroup straw poll.</li> <li>Incorporated workgroup feedback.</li> </ul>
2019-06-12	<ul style="list-style-type: none"> <li>Added full version of Figure 247: Identify – Identify Controller Data Structure.</li> <li>Added scope column to Figure 247 to specify the scope of a field in the Identify Controller Data structure.</li> <li>Updated field in Figure 247 to make them generic to all I/O Command Sets.</li> </ul>
2019-06-19	<ul style="list-style-type: none"> <li>Minor updates from workgroup feedback.</li> <li>Changed Scope to Property in Identify Controller data structure and added descriptive footnote.</li> </ul>
2019-06-23	<ul style="list-style-type: none"> <li>Updated Property field in Identify Controller Data Structure to map items to the proper property</li> <li>Removed comments associated with review feedback that was discussed in the workgroup and resolved.</li> </ul>
2019-06-26	<ul style="list-style-type: none"> <li>Updated TP based on workgroup feedback.</li> </ul>
2019-07-06	<ul style="list-style-type: none"> <li>Updated TP based on workgroup feedback.</li> <li>Fixed cross references and Figure numbering.</li> </ul>
2019-07-14	<ul style="list-style-type: none"> <li>Updated TP based on workgroup feedback</li> <li>Added Number of Simultaneous Command Set Combinations to the Identify Command Set Data Structure to that the host knows the number of entries to expect without scanning for a value of zero.</li> <li>Changed name of Command Set Not Supported in Command Set Profile to Command Set Not Enabled.</li> <li>Incorporated NVMe-MI changes proposed by Mike Allison.</li> </ul>
2019-07-15	<ul style="list-style-type: none"> <li>Incorporated review feedback from Mike Allison. <ul style="list-style-type: none"> <li>Updated the Number of Simultaneous Command Set Combinations field to indicate that it is a zeros based value and that values greater than 1020 are reserved.</li> <li>Updated Controller Identifier (CTRLID) field to include Optional Command Supported and list Data Structure Type (DTYP) field.</li> </ul> </li> </ul>
2019-07-23	<ul style="list-style-type: none"> <li>Changed the name of Identify Command Set to Identify I/O Command Set</li> <li>Changed the name of the Identify Command Set Data Structure to Identify I/O Command Set Data Structure.</li> <li>Reworked sections 5.15.2 and 5.21.1.TBDp to align with the direction agreed to in the 7/18 workgroup meeting. There is a list of supported command set vectors and we will use a Feature value to select which one is used.</li> <li>Updated section 5.15.2.3 and 5.15.2.6 so that NST is used for filtering and not checking.</li> <li>Updated text to incorporate NVMe-MI related feedback from Austin Bolen.</li> </ul>
2019-10-01	<ul style="list-style-type: none"> <li>Added sections for multiple command sets: <ul style="list-style-type: none"> <li>1.6.24</li> <li>1.6.37</li> <li>4.4</li> <li>Figure 118</li> <li>Figure 130</li> <li>Figure TBDXX</li> </ul> </li> </ul>
2019-10-03	<ul style="list-style-type: none"> <li>Results of Technical Work group meeting</li> </ul>

10-16-2019	<ul style="list-style-type: none"> <li>• Evaluation of occurrences of NVM Command and NVM Command Set through Get Log Page</li> </ul>
10-24-2019	<ul style="list-style-type: none"> <li>• Moved “user data” and “logical block data” changes to ECN</li> <li>• Updated Submission Queue and Completion Queue entries</li> </ul>
11-04-2019	<ul style="list-style-type: none"> <li>• Continued adding sections for removal/modification of NVM Command Set through the beginning of 5.21</li> </ul>
11-05-2019	<ul style="list-style-type: none"> <li>• Combined Status tables and added command set columns</li> <li>•</li> </ul>
11-13-2019	<ul style="list-style-type: none"> <li>• Added column to table 184</li> <li>• Added figure 197</li> <li>• Added figure 207</li> <li>• Added section 5.21.1</li> <li>• Added Format command</li> </ul>
11-20-2019	<ul style="list-style-type: none"> <li>• Added CQE dword 1 to section 5 and 6</li> <li>• Added command table modifications to section 6 for refactoring</li> <li>• Updated the Format command</li> </ul>
12-03-2019	<ul style="list-style-type: none"> <li>• Removed Namespace Type and made all qualifications/references to I/O command set</li> <li>• Change tracking turned on to easily find changes.</li> </ul>
12-05-2019	<ul style="list-style-type: none"> <li>• Modified Identify Controller to look at CSI field for specific data structure to return</li> </ul>
12-09-2019	<ul style="list-style-type: none"> <li>• Made I/O Command Set consistent</li> <li>• Incorporated comments from Mike Allison and Christoph Hellwig per file sent in email to reflector.</li> </ul>
12-10-2019 rev a	<ul style="list-style-type: none"> <li>• Additional comments from Christoph incorporated</li> </ul>
12-11-2019	<ul style="list-style-type: none"> <li>• Removed “Conventional Logical Block namespace” and only reference the NVM Command Set</li> </ul>
12-12-2019	<ul style="list-style-type: none"> <li>• Removed change tracking that was discussed in technical work group</li> <li>• Changed response to Identify controller for Command Sets that do not require a data structure for this to return a zero filled data structure</li> <li>• Ready for phase 2 exit</li> </ul>
01-14-2020	<ul style="list-style-type: none"> <li>• Minor updates from Christoph</li> <li>• Accepted changes reviewed on 12/12/2019</li> <li>•</li> </ul>
01-15-2020	<ul style="list-style-type: none"> <li>• Modified TBDp1 to reflect 9 bits to specify the command set combination</li> <li>• Modified TBDa: Identify I/O Command Set Data Structure, to reflect that there are 510 combinations not 509.</li> </ul>
01-16-2020	<ul style="list-style-type: none"> <li>• Incorporate comments from Mike Allison</li> </ul>
02-05-2020	<ul style="list-style-type: none"> <li>• Incorporated agreed changes from 01-16-2020 meeting</li> <li>• Removed comments that were resolved prior to 01-16-2020 meeting</li> <li>• Incorporated additional changes from Christoph</li> </ul>
02-12-2020	<ul style="list-style-type: none"> <li>• Removed all resolved comments</li> <li>• Removed change tracking from previously reviewed modifications</li> <li>• Incorporated changes for reservations</li> <li>• Incorporated changes for the wording of read commands and write commands</li> <li>• Added section 2.3.2 from merged document from the Refactoring group</li> </ul>
2-13-2020	<ul style="list-style-type: none"> <li>• Removed all resolved comments</li> <li>• Removed change tracking for changes previously reviewed</li> <li>• Added modifications from 2-13-2020 technical workgroup meeting</li> </ul>
2-18-2020	<ul style="list-style-type: none"> <li>• Removed all resolved comments and change modifications from previous revisions</li> <li>• Made editorial changes from Mike Allison</li> <li>• Made editorial changes from Christoph Helwig</li> <li>• Added column to Figure 139: Opcodes for Admin Commands</li> </ul>

2-21-2020	<ul style="list-style-type: none"> <li>Fixed section and figure numbering</li> <li>Changed I/O Command Set Not Supported to be applicable for Namespace Management as well</li> <li>Fixed multiple occurrences of “active namespaces” to “active NSIDs”</li> <li>Delete 6.NVMTBD1 Host Read Commands and Host Write Commands</li> <li>Editorial fixes</li> <li>Added Identify Controller data structure for NVM Command Set</li> </ul>
2-26-2020	<ul style="list-style-type: none"> <li>Removed all previous change tracking</li> <li>Added FIVE new CNS values to make backwards compatibility for NVM command set controllers and hosts.</li> <li>Added comment on Format NVM command that needs to be addressed</li> </ul>
2-27-2020	<ul style="list-style-type: none"> <li>Added comments from David Black and Christoph</li> </ul>
3-4-2020	<ul style="list-style-type: none"> <li>Added modifications from 2-27 conference call</li> <li>Added bit 3 to the FNA field in Identify and Format NVM scope</li> </ul>
3-5-2020	<ul style="list-style-type: none"> <li>Added qualifiers on FNA bits 0 and 1 that if bit 3 is set to ‘1’, these bits are cleared to ‘0’</li> </ul>
3-12-2020	<ul style="list-style-type: none"> <li>Added Command Specific Dword 2 and Command Specific Dword 3 to the Common command data structure</li> <li>Removed restrictions on the return of CNS unless implementing a specific I/O Command Set</li> <li>Fixed 2 places that were not fixed when going from CSI value 1h to CSI value 0h</li> </ul>
4-29-2020	<ul style="list-style-type: none"> <li>Modifications for member review comments</li> <li>Used the Log Specific Identifier in the Get Features command for Command Set Identifier</li> </ul>
4-30-2020	<ul style="list-style-type: none"> <li>Removed completed comments</li> <li>Moved Command Set Identifier from Log Specific Identifier field into a new field in CDW-14</li> </ul>
5-7-2020	<ul style="list-style-type: none"> <li>Added wording in AEN handling (from technical call)</li> <li>Added wording for handling of Set Features command for the I/O Command Set Profile if CC.CSS is set to any value other than 110b</li> </ul>
06-07-2020	<ul style="list-style-type: none"> <li>Integrated into the NVM Express Base Specification and NVMe MI Specification.</li> </ul>
06-11-2020	<ul style="list-style-type: none"> <li>Integration corrections</li> </ul>
06-23-2020	<ul style="list-style-type: none"> <li>Updated value for Set Features status code for I/O Command Set Combination Rejected.</li> </ul>
06-26-2020	<ul style="list-style-type: none"> <li>Updated figure 265 (Namespace Management command) to have the Identify Namespace field be defined by section 5.20.x.</li> </ul>
07-09-2020	<ul style="list-style-type: none"> <li>Added formatted into definition of namespace in three places</li> </ul>
07-22-2020	<ul style="list-style-type: none"> <li>This replaces the 4056a that is in member review. 4056a will not become published as this ECN fixes known problems in the member review version.</li> <li>Modified Stats code returned for CNS 16h to match the code for CNS 00h and 11h</li> <li>In Identify Controller data structure, I/O Command Set independent, ONCS specified Dataset Management as an NVM Command Set command. removed this phrase from 2 places in this field.</li> <li></li> </ul>
07-28-2020	<ul style="list-style-type: none"> <li>Reversed change in ONCS, made the reference in figure 346 point to the NVM command set added footnote 3 on Dataset Management in figure 462</li> <li>Made Format in progress status code a command set specific status code</li> <li>Added IO Command Set Profile to figure 422</li> </ul>

07-30-2020	<ul style="list-style-type: none"> <li>Modified the name of 5.15.2.TBDa.1 and the figure in that section</li> <li>Put section numbers back into table 346 and removed key for NVM Command Set</li> </ul>
09-09-2020	<ul style="list-style-type: none"> <li>Added "Namespace Management" to one row in figure 128 as per member review comment.</li> </ul>
09-15-2020	<ul style="list-style-type: none"> <li>Integrated into the NVMe Management Interface Specification, Revision 1.1 and the NVMe Express Base Specification.</li> </ul>
09-16-2020	<ul style="list-style-type: none"> <li>Accepted all changes for ratification</li> </ul>
10-14-2020	<ul style="list-style-type: none"> <li>Fixed the Namespace Identification Descriptor list to indicate that a NIDT=4 is required if one or more I/O Command Sets are supported.</li> </ul>
10-15-2020	<ul style="list-style-type: none"> <li>Requirement changed from 10-14-2020 change to just state that if CAP.CSS bit 43 is set to a value to be consistent with exiting usage in the specification.</li> </ul>
10-30-2020	<ul style="list-style-type: none"> <li>Added section 7.1.2 to show the support requirements for the I/O Command Set Profile feature for an administrative controller.</li> </ul>
11-03-2020	<ul style="list-style-type: none"> <li>Added section 5.14.1.4 which is moved from TP 4056a discussion.</li> </ul>
11-5-2020	<ul style="list-style-type: none"> <li>Accepted all changes for a review.</li> </ul>
12-2-2020	<ul style="list-style-type: none"> <li>Modified Namespace Attribute Changed field in Figure 147 (Asynchronous Event Information – Notice) to include changes to any supported Identify I/O Command Set Specific Namespace data structure.</li> </ul>
12-9-2020	<ul style="list-style-type: none"> <li>Renamed TP 4056d to follow policy that lower letter revision change only changes after a ratification.</li> </ul>
12-10-2020	<ul style="list-style-type: none"> <li>Accepted all changes for member review.</li> </ul>
1-28-2021	<ul style="list-style-type: none"> <li>Fixed references to CAP register instead of CAP.CSS. Fixed name of references external specification. New text introduced by this TP uses the phrase "a status code of xxxx". Removed comments.</li> </ul>
2-2-2021	<ul style="list-style-type: none"> <li>Integrated into the NVMe Base Specification and the NVMe Management Interface Specification, Revision 1.1.</li> </ul>
2-3-2021	<ul style="list-style-type: none"> <li>Accepted all changes and removed comments.</li> </ul>
2-4-2021	<ul style="list-style-type: none"> <li>Changed "status of code" to "status code of"</li> </ul>

## Description of Specification Changes

Markup Conventions:

Black:	Unchanged (however, hot links are removed)
<del>Red Strikethrough:</del>	Deleted
Red:	New
Red Highlighted:	TBD values, anchors, and links to be inserted.
<Green Bracketed>:	Notes to editor

## Description for NVMe 1.4 Changes Document

This technical proposal adds support for namespace types. NVMe currently only supports namespaces that contain traditional logical blocks that may be randomly read and written. In some applications there is value in exposing namespaces that contain logical blocks that have special access rules (e.g., must be written sequentially), namespaces that contain logical blocks that are used to perform data transformations, or namespaces that contain things other than logical blocks. The purpose of this technical proposal is to extend the NVMe architecture to support different types of namespaces. This extension is fully backward compatible with the existing NVMe specification.

## NVM Express specification 1.4 Changes

### Modify Section 1.4 Theory of Operation as shown below:

The interface has the following key attributes:

- Does not require uncacheable / MMIO register reads in the command submission or completion path;
- A maximum of one MMIO register write is necessary in the command submission path;
- Support for up to 65,535 I/O Queues, with each I/O Queue supporting up to 65,535 outstanding commands;
- Priority associated with each I/O Queue with well-defined arbitration mechanism;
- All information to complete a 4 KiB read request is included in the 64B command itself, ensuring efficient small I/O operation;
- Efficient and streamlined command sets;
- Support for MSI/MSI-X and interrupt aggregation;
- Support for multiple namespaces;
- Efficient support for I/O virtualization architectures like SR-IOV;
- Robust error reporting and management capabilities; and
- Support for multi-path I/O and namespace sharing.

This specification defines a streamlined set of registers whose functionality includes:

- Indication of controller capabilities;
- Status for controller failures (command status is processed via CQ directly);
- Admin Queue configuration (I/O Queue configuration processed via Admin commands); and
- Doorbell registers for scalable number of Submission and Completion Queues.

An NVM Express controller is associated with a single PCI Function. The capabilities and settings that apply to the entire controller are indicated in the Controller Capabilities (CAP) register and the Identify Controller data structure.

A namespace is a ~~formatted~~ quantity of non-volatile memory ~~that may be accessed by a host-formatted into logical blocks~~. Associated with each namespace is an I/O Command Set that operates on that namespace.

An NVM Express controller may support multiple namespaces that are referenced using a namespace ID. Namespaces may be created and deleted using the Namespace Management and Namespace Attachment commands. The Identify Namespace data structure ~~and the Identify I/O Command Set Specific Namespace data structure~~ indicates capabilities and settings that are specific to a particular namespace. ~~The capabilities and settings that are common to all namespaces are reported by the Identify Namespace data structure for namespace ID FFFFFFFFh.~~

The NVM Express interface is based on a paired Submission and Completion Queue mechanism. Commands are placed by host software into a Submission Queue. Completions are placed into the associated Completion Queue by the controller. Multiple Submission Queues may utilize the same Completion Queue. Submission and Completion Queues are allocated in memory.

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

An I/O Command Set is used with an I/O queue pair. This specification defines one I/O Command Set, named the NVM Command Set. ~~Additional I/O Command Sets are defined in other specifications (e.g., Key Value Command Set, or Zoned Namespace Command Set). The host selects one I/O Command Set that is used for all I/O queue pairs.~~

## Add subsection after Section 1.4.1 Multi-Path I/O and Namespace Sharing

### 1.4.~~TBD~~ Command Sets

A controller may simultaneously support multiple I/O Command Sets. The I/O Command Sets that a controller supports and which of these I/O Command Sets the controller simultaneously supports is reported in the Identify I/O Command Set data structure (refer to section ~~5.15.2.TBD~~). The contents of the Identify I/O Command Set data structure are not required to be the same for all controllers in an NVM subsystem.

A namespace is associated with exactly one I/O Command Set. For I/O commands and I/O Command Set Specific Admin commands, the I/O Command Set with which a Submission Queue entry is associated is determined by the Namespace Identifier (NSID) field in the command.

An NVM subsystem may contain namespaces each of which is associated with a different I/O Command Set. A controller may support attached namespaces that use any of the I/O Command Sets that the controller simultaneously supports as indicated in the I/O Command Set Profile (refer to section ~~5.21.1.TBD~~).

## Modify Section 1.6 Definitions as shown below and add new entries in alphabetical order:

### 1.6.25 namespace

A ~~formatted~~ quantity of non-volatile memory that may be directly accessed by a host. ~~formatted into logical blocks. When formatted, a namespace of size n is a collection of logical blocks with logical block addresses from 0 to (n-1) directly accessible by a host.~~



Modify Section 3.1.1 as shown below:

### 3.1.1 Offset 0h: CAP – Controller Capabilities

This register indicates basic capabilities of the controller to host software.

Figure 69: Offset 0h: CAP – Controller Capabilities

Bits	Type	Reset	Description																	
63: 58	RO	0h	Reserved																	
57	RO	Impl Spec	<b>Controller Memory Buffer Supported (CMBS):</b> If set to ‘1’, this bit indicates that the controller supports the Controller Memory Buffer, and that addresses supplied by the host are permitted to reference the Controller Memory Buffer only if the host has enabled the Controller Memory Buffer’s controller memory space.  If the controller supports the Controller Memory Buffer, this bit shall be set to ‘1’.																	
...	...	...	...																	
45	RO	Impl Spec	<b>Boot Partition Support (BPS):</b> This bit indicates whether the controller supports Boot Partitions. If this bit is set to ‘1’, the controller supports Boot Partitions. If this bit is cleared to ‘0’, the controller does not support Boot Partitions. Refer to section 8.13.																	
44:37	RO	Impl Spec	<b>Command Sets Supported (CSS):</b> This field indicates the I/O Command Set(s) that the controller supports.  <del>If a bit is set to ‘1’, then the corresponding I/O Command Set. If a bit is cleared to ‘0’, then the corresponding I/O Command Set is not supported.</del>  Bit 44 is set to ‘1’ if no I/O Command Set is supported.																	
			Bits	Definition	37	<del>NVM cCommand sSet is supported</del> Controller supports the NVM Command Set. Controllers that support the NVM Command Set shall set this bit even if bit 43 is set to ‘1’.	42:38	Reserved	39	Reserved	40	Reserved	41	Reserved	42	Reserved	43	<del>Reserved</del> Controller supports one or more I/O Command Sets and supports the Identify I/O Command Set data structure (refer to section 5.15.2.TBD). Controllers that support I/O Command Sets other than the NVM Command Set shall set bit 43 to ‘1’. Controllers that only support the NVM Command Set may set this bit to ‘1’ to indicate support for the Command Set Identifier field in commands that use the Command Set Identifier field.	44	No I/O Command Set is supported (i.e., only the Admin Command Set is supported)
			Bits	Definition																
			37	<del>NVM cCommand sSet is supported</del> Controller supports the NVM Command Set. Controllers that support the NVM Command Set shall set this bit even if bit 43 is set to ‘1’.																
			42:38	Reserved																
			39	Reserved																
			40	Reserved																
			41	Reserved																
			42	Reserved																
			43	<del>Reserved</del> Controller supports one or more I/O Command Sets and supports the Identify I/O Command Set data structure (refer to section 5.15.2.TBD). Controllers that support I/O Command Sets other than the NVM Command Set shall set bit 43 to ‘1’. Controllers that only support the NVM Command Set may set this bit to ‘1’ to indicate support for the Command Set Identifier field in commands that use the Command Set Identifier field.																
44	No I/O Command Set is supported (i.e., only the Admin Command Set is supported)																			

**Figure 69: Offset 0h: CAP – Controller Capabilities**

Bits	Type	Reset	Description
36	RO	Impl Spec	<b>NVM Subsystem Reset Supported (NSSRS):</b> This bit indicates whether the controller supports the NVM Subsystem Reset feature defined in 7.3.1. This bit is set to '1' if the controller supports the NVM Subsystem Reset feature. This bit is cleared to '0' if the controller does not support the NVM Subsystem Reset feature.
...	...	...	...
15:00	RO	Impl Spec	<b>Maximum Queue Entries Supported (MQES):</b> This field indicates the maximum individual queue size that the controller supports. For NVMe over PCIe implementations, this value applies to the I/O Submission Queues and I/O Completion Queues that the host creates. For NVMe over Fabrics implementations, this value applies to only the I/O Submission Queues that the host creates. This is a 0's based value. The minimum value is 1h, indicating two entries.

Modify Section 3.1.5 as shown below:

### 3.1.5 Offset 14h: CC – Controller Configuration

This register modifies settings for the controller. Host software shall set the Arbitration Mechanism (CC.AMS), the Memory Page Size (CC.MPS), and the Command Set (CC.CSS) to valid values prior to enabling the controller by setting CC.EN to '1'. Attempting to create an I/O queue before initializing the I/O Completion Queue Entry Size (CC.IOCQES) and the I/O Submission Queue Entry Size (CC.IOSQES) should cause a controller to abort a Create I/O Completion Queue command or a Create I/O Submission Queue command with a status code of Invalid Queue Size.

**Figure 78: Offset 14h: CC – Controller Configuration**

Bits	Type	Reset	Description
31:24	RO	0h	Reserved
...	...	...	...
10:07	RW	0h	<b>Memory Page Size (MPS):</b> This field indicates the host memory page size. The memory page size is $(2^{(12 + MPS)})$ . Thus, the minimum host memory page size is 4 KiB and the maximum host memory page size is 128 MiB. The value set by host software shall be a supported value as indicated by the CAP.MPSMAX and CAP.MPSMIN fields. This field describes the value used for PRP entry size. This field shall only be modified when EN is cleared to '0'.

Figure 78: Offset 14h: CC – Controller Configuration

Bits	Type	Reset	Description										
06:04	RW	000b	<p><b>I/O Command Set Selected (CSS):</b> This field specifies the I/O Command Set that is selected. Host software shall only select a supported I/O Command Set, as indicated in CAP.CSS. This field shall only be changed when the controller is disabled (CC.EN is cleared to '0'). The I/O Command Set selected shall be used for all I/O Submission Queues. This field specifies the I/O Command Set or Sets that are selected. This field shall only be changed when the controller is disabled (CC.EN is cleared to '0'). The I/O Command Set or Sets that are selected shall be used for all I/O Submission Queues.</p> <p>If bit 37 is set to '1' in the CAP register, then a value of 000b selects the NVM Command Set. If bit 37 is set to '0' in the CAP register then the value 000b is reserved.</p> <p>If bit 43 is set to '1' in the CAP register, then a value of 110b selects all I/O Command Sets supported by the controller. The I/O Command Sets that are supported in this case are reported in the Identify I/O Command Set data structure (refer to section 5.15.2.TBD). If bit 43 is set to '0' in the CAP register, then the value 110b is reserved.</p> <p>If bit 44 is set to '1' in the CAP register, then the value 111b indicates that only the Admin Command Set is supported and that no I/O Command Set or I/O Command Set Specific Admin commands are supported. When only the Admin Command Set is supported, any command submitted on an I/O Submission Queue and any I/O Command Set Specific Admin command submitted on the Admin Submission Queue is completed with status Invalid Command Opcode. If bit 44 is cleared to '0' in the CAP register, then the value of 111b is reserved.</p> <table><tr><th>Value</th><th>Definition</th></tr><tr><td>000b</td><td>NVM Command Set</td></tr><tr><td>001b to 101b</td><td>Reserved</td></tr><tr><td>110b</td><td>All supported I/O Command Sets</td></tr><tr><td>111b</td><td>Admin Command Set only</td></tr></table>	Value	Definition	000b	NVM Command Set	001b to 101b	Reserved	110b	All supported I/O Command Sets	111b	Admin Command Set only
Value	Definition												
000b	NVM Command Set												
001b to 101b	Reserved												
110b	All supported I/O Command Sets												
111b	Admin Command Set only												
03:01	RO	000b	Reserved										
...	...	...	...										

Add new section 4.TBD as shown below:

#### 4.TBD I/O Command Sets

I/O commands perform operations on namespaces, and each namespace is associated with exactly one I/O Command Set. For example, commands in the NVM Command Set access data represented in a namespace as logical blocks, and commands in the Key-Value Command set access data represented in a namespace as key-value pairs.

The association of a namespace to an I/O Command Set is specified when the namespace is created and is fixed for the lifetime of the namespace.

A controller may support one or more I/O Command Sets. If the controller supports more than one I/O Command Set, then the controller may have namespaces attached that are associated with I/O Command Sets as described in 5.21.1.TBDp. I/O commands are processed based on the I/O Command Set associated with the namespace that processes the command.

**Modify Section 4.2 Submission Queue Entry – Command Format as shown below:**

## 4.2 Submission Queue Entry – **Common** Command Format

Each **Common Command Format** command is 64 bytes in size.

Command Dword 0, Namespace Identifier, Metadata Pointer, PRP Entry 1, PRP Entry 2, SGL Entry 1, and Metadata SGL Segment Pointer have common definitions for all Admin commands and **NVM I/O** commands **for all I/O Command Sets**. Metadata Pointer, PRP Entry 1, PRP Entry 2, and Metadata SGL Segment Pointer are not used by all commands. Command Dword 0 is defined in Figure 104.

**Figure 104: Command Dword 0**

Bits	Description										
31:16	<b>Command Identifier (CID):</b> This field specifies a unique identifier for the command when combined with the Submission Queue identifier.										
15:14	<p><b>PRP or SGL for Data Transfer (PSDT):</b> This field specifies whether PRPs or SGLs are used for any data transfer associated with the command. PRPs shall be used for all Admin commands for NVMe over PCIe implementations. SGLs shall be used for all Admin and I/O commands for NVMe over Fabrics implementations. This field shall be set to 01b for NVMe over Fabrics revision 1.0 implementations. The definition is described in the table below.</p> <table border="1"> <thead> <tr> <th>Value</th><th>Definition</th></tr> </thead> <tbody> <tr> <td>00b</td><td>PRPs are used for this transfer.</td></tr> <tr> <td>01b</td><td>SGLs are used for this transfer. If used, Metadata Pointer (MPTR) contains an address of a single contiguous physical buffer that is byte aligned.</td></tr> <tr> <td>10b</td><td>SGLs are used for this transfer. If used, Metadata Pointer (MPTR) contains an address of an SGL segment containing exactly one SGL Descriptor that is qword aligned.</td></tr> <tr> <td>11b</td><td>Reserved</td></tr> </tbody> </table> <p>If there is metadata that is not interleaved with the logical block data, as specified in the Format NVM command, then the Metadata Pointer (MPTR) field is used to point to the metadata. The definition of the Metadata Pointer field is dependent on the setting in this field. Refer to Figure 105.</p>	Value	Definition	00b	PRPs are used for this transfer.	01b	SGLs are used for this transfer. If used, Metadata Pointer (MPTR) contains an address of a single contiguous physical buffer that is byte aligned.	10b	SGLs are used for this transfer. If used, Metadata Pointer (MPTR) contains an address of an SGL segment containing exactly one SGL Descriptor that is qword aligned.	11b	Reserved
Value	Definition										
00b	PRPs are used for this transfer.										
01b	SGLs are used for this transfer. If used, Metadata Pointer (MPTR) contains an address of a single contiguous physical buffer that is byte aligned.										
10b	SGLs are used for this transfer. If used, Metadata Pointer (MPTR) contains an address of an SGL segment containing exactly one SGL Descriptor that is qword aligned.										
11b	Reserved										
13:10	Reserved										

**Figure 104: Command Dword 0**

Bits	Description										
09:08	<p><b>Fused Operation (FUSE):</b> In a fused operation, a complex command is created by “fusing” together two simpler commands. Refer to section 4.12. This field specifies whether this command is part of a fused operation and if so, which command it is in the sequence.</p> <table> <tr> <th>Value</th><th>Definition</th></tr> <tr> <td>00b</td><td>Normal operation</td></tr> <tr> <td>01b</td><td>Fused operation, first command</td></tr> <tr> <td>10b</td><td>Fused operation, second command</td></tr> <tr> <td>11b</td><td>Reserved</td></tr> </table>	Value	Definition	00b	Normal operation	01b	Fused operation, first command	10b	Fused operation, second command	11b	Reserved
Value	Definition										
00b	Normal operation										
01b	Fused operation, first command										
10b	Fused operation, second command										
11b	Reserved										
07:00	<b>Opcode (OPC):</b> This field specifies the opcode of the command to be executed.										

The ~~64 byte Common~~ Command Format ~~for the Admin Command Set and NVM Command Set~~ is defined in Figure 105.

SGLs shall not be used for Admin commands in NVMe over PCIe implementations.

**Figure 105: Command Format ~~Admin and NVM Command Set~~ for the Common Command Format**

Bytes	Description
03:00	<b>Command Dword 0 (CDW0):</b> This field is common to all commands and is defined in Figure 104.
07:04	<p><b>Namespace Identifier (NSID):</b> This field specifies the namespace that this command applies to. If the namespace identifier is not used for the command, then this field shall be cleared to 0h. The value FFFFFFFFh in this field is a broadcast value (refer to section 6.1, where the scope (e.g., the NVM subsystem, all attached namespaces, or all namespaces in the NVM subsystem) is dependent on the command. Refer to Figure 139, Figure 140, and Figure 246 for commands that support the use of the value FFFFFFFFh in this field.</p> <p>Specifying an inactive namespace identifier (refer to section 6.1.4) in a command that uses the namespace identifier shall cause the controller to abort the command with status Invalid Field in Command, unless otherwise specified. Specifying an invalid namespace identifier (refer to section 6.1.2 in a command that uses the namespace identifier shall cause the controller to abort the command with status Invalid Namespace or Format, unless otherwise specified.</p> <p>If the namespace identifier is used for the command (refer to Figure 139 and Figure 140), the value FFFFFFFFh is not supported for that command, and the host specifies a value of FFFFFFFFh, then the controller should abort the command with status Invalid Field in Command, unless otherwise specified.</p> <p>If the namespace identifier is not used for the command and the host specifies a value from 1h to FFFFFFFEh, then the controller should abort the command with status Invalid Field in Command, unless otherwise specified.</p>
<del>115:08</del>	<del>Reserved</del> <b>Command Dword 2 (CDW2):</b> This field is command specific Dword 2
<del>15:12</del>	<del>Command Dword 3 (CDW3):</del> This field is command specific Dword 3

**Figure 105: Command Format ~~Admin and NVM Command Set for the Common Command Format~~**

Bytes	Description
23:16	<p><b>Metadata Pointer (MPTR):</b> This field is valid only if the command has metadata that is not interleaved with the logical block data, as specified in the Format NVM command. This is a reserved field in NVMe over Fabrics implementations.</p> <p>If CDW0.PSDT (refer to Figure 104) is cleared to 00b, then this field shall contain the address of a contiguous physical buffer of metadata and that address shall be dword aligned (i.e., bits 1:0 cleared to 00b). The controller is not required to check that bits 1:0 are cleared to 00b. The controller may report an error of Invalid Field in Command if bits 1:0 are not cleared to 00b. If the controller does not report an error of Invalid Field in Command, then the controller shall operate as if bits 1:0 are cleared to 00b.</p> <p>If CDW0.PSDT is set to 01b, then this field shall contain the address of a contiguous physical buffer of metadata and that address may be aligned on any byte boundary.</p> <p>If CDW0.PSDT is set to 10b, then this field shall contain the address of an SGL segment that contains exactly one SGL Descriptor. The address of that SGL segment shall be qword aligned (i.e., bits 2:0 cleared to 000b). The SGL Descriptor contained in that SGL segment is the first SGL Descriptor of the metadata for the command. If the SGL Descriptor contained in that SGL segment is an SGL Data Block descriptor, then that SGL Data Block Descriptor is the only SGL Descriptor and therefore describes the entire metadata data transfer. Refer to section 4.4. The controller is not required to check that bits 2:0 are cleared to 000b. The controller may report an error of Invalid Field in Command if bits 2:0 are not cleared to 000b. If the controller does not report an error of Invalid Field in Command, then the controller shall operate as if bits 2:0 are cleared to 000b.</p>

**Figure 105: Command Format ~~Admin and NVM Command Set~~ for the Common Command Format**

Bytes	Description						
39:24	<p><b>Data Pointer (DPTR):</b> This field specifies the data used in the command.</p> <p>If CDW0.PSDT is cleared to 00b, then the definition of this field is:</p> <table border="1"> <tr> <td>39:32</td><td> <p><b>PRP Entry 2 (PRP2):</b> This field:</p> <ul style="list-style-type: none"> <li>a) is reserved if the data transfer does not cross a memory page boundary;</li> <li>b) specifies the Page Base Address of the second memory page if the data transfer crosses exactly one memory page boundary. E.g.,: <ul style="list-style-type: none"> <li>i. the command data transfer length is equal in size to one memory page and the offset portion of the PBAO field of PRP1 is non-zero; or</li> <li>ii. the Offset portion of the PBAO field of PRP1 is equal to 0h and the command data transfer length is greater than one memory page and less than or equal to two memory pages in size;</li> </ul> </li> <li>• and</li> <li>c) is a PRP List pointer if the data transfer crosses more than one memory page boundary. E.g.,: <ul style="list-style-type: none"> <li>i. the command data transfer length is greater than or equal to two memory pages in size but the offset portion of the PBAO field of PRP1 is non-zero; or</li> <li>ii. the command data transfer length is equal in size to more than two memory pages and the Offset portion of the PBAO field of PRP1 is equal to 0h.</li> </ul> </li> </ul> </td></tr> <tr> <td>31:24</td><td> <p><b>PRP Entry 1 (PRP1):</b> This field contains the first PRP entry for the command or a PRP List pointer depending on the command.</p> </td></tr> </table> <p>If CDW0.PSDT is set to 01b or 10b, then the definition of this field is:</p> <table border="1"> <tr> <td>39:24</td><td> <p><b>SGL Entry 1 (SGL1):</b> This field contains the first SGL segment for the command. If the SGL segment is an SGL Data Block or Keyed SGL Data Block or Transport SGL Data Block descriptor, then it describes the entire data transfer. If more than one SGL segment is needed to describe the data transfer, then the first SGL segment is a Segment, or Last Segment descriptor. Refer to section 4.4 for the definition of SGL segments and descriptor types.</p> <p>The NVMe Transport may support a subset of SGL Descriptor types and features as defined in the NVMe Transport binding specification.</p> </td></tr> </table>	39:32	<p><b>PRP Entry 2 (PRP2):</b> This field:</p> <ul style="list-style-type: none"> <li>a) is reserved if the data transfer does not cross a memory page boundary;</li> <li>b) specifies the Page Base Address of the second memory page if the data transfer crosses exactly one memory page boundary. E.g.,: <ul style="list-style-type: none"> <li>i. the command data transfer length is equal in size to one memory page and the offset portion of the PBAO field of PRP1 is non-zero; or</li> <li>ii. the Offset portion of the PBAO field of PRP1 is equal to 0h and the command data transfer length is greater than one memory page and less than or equal to two memory pages in size;</li> </ul> </li> <li>• and</li> <li>c) is a PRP List pointer if the data transfer crosses more than one memory page boundary. E.g.,: <ul style="list-style-type: none"> <li>i. the command data transfer length is greater than or equal to two memory pages in size but the offset portion of the PBAO field of PRP1 is non-zero; or</li> <li>ii. the command data transfer length is equal in size to more than two memory pages and the Offset portion of the PBAO field of PRP1 is equal to 0h.</li> </ul> </li> </ul>	31:24	<p><b>PRP Entry 1 (PRP1):</b> This field contains the first PRP entry for the command or a PRP List pointer depending on the command.</p>	39:24	<p><b>SGL Entry 1 (SGL1):</b> This field contains the first SGL segment for the command. If the SGL segment is an SGL Data Block or Keyed SGL Data Block or Transport SGL Data Block descriptor, then it describes the entire data transfer. If more than one SGL segment is needed to describe the data transfer, then the first SGL segment is a Segment, or Last Segment descriptor. Refer to section 4.4 for the definition of SGL segments and descriptor types.</p> <p>The NVMe Transport may support a subset of SGL Descriptor types and features as defined in the NVMe Transport binding specification.</p>
39:32	<p><b>PRP Entry 2 (PRP2):</b> This field:</p> <ul style="list-style-type: none"> <li>a) is reserved if the data transfer does not cross a memory page boundary;</li> <li>b) specifies the Page Base Address of the second memory page if the data transfer crosses exactly one memory page boundary. E.g.,: <ul style="list-style-type: none"> <li>i. the command data transfer length is equal in size to one memory page and the offset portion of the PBAO field of PRP1 is non-zero; or</li> <li>ii. the Offset portion of the PBAO field of PRP1 is equal to 0h and the command data transfer length is greater than one memory page and less than or equal to two memory pages in size;</li> </ul> </li> <li>• and</li> <li>c) is a PRP List pointer if the data transfer crosses more than one memory page boundary. E.g.,: <ul style="list-style-type: none"> <li>i. the command data transfer length is greater than or equal to two memory pages in size but the offset portion of the PBAO field of PRP1 is non-zero; or</li> <li>ii. the command data transfer length is equal in size to more than two memory pages and the Offset portion of the PBAO field of PRP1 is equal to 0h.</li> </ul> </li> </ul>						
31:24	<p><b>PRP Entry 1 (PRP1):</b> This field contains the first PRP entry for the command or a PRP List pointer depending on the command.</p>						
39:24	<p><b>SGL Entry 1 (SGL1):</b> This field contains the first SGL segment for the command. If the SGL segment is an SGL Data Block or Keyed SGL Data Block or Transport SGL Data Block descriptor, then it describes the entire data transfer. If more than one SGL segment is needed to describe the data transfer, then the first SGL segment is a Segment, or Last Segment descriptor. Refer to section 4.4 for the definition of SGL segments and descriptor types.</p> <p>The NVMe Transport may support a subset of SGL Descriptor types and features as defined in the NVMe Transport binding specification.</p>						
43:40	<b>Command Dword 10 (CDW10):</b> This field is command specific Dword 10.						
47:44	<b>Command Dword 11 (CDW11):</b> This field is command specific Dword 11.						
51:48	<b>Command Dword 12 (CDW12):</b> This field is command specific Dword 12.						
55:52	<b>Command Dword 13 (CDW13):</b> This field is command specific Dword 13.						
59:56	<b>Command Dword 14 (CDW14):</b> This field is command specific Dword 14.						
63:60	<b>Command Dword 15 (CDW15):</b> This field is command specific Dword 15.						

In addition to the fields commonly defined for the ~~Common Command Format~~ ~~all Admin and NVM commands~~, Admin and NVM Vendor Specific commands may support the Number of Dwords in Data Transfer and Number of Dwords in Metadata Transfer fields. If supported, the command format for the

Admin Vendor Specific Command and NVM Vendor Specific Commands are defined in Figure 106. For more details, refer to section 8.7.

**Figure 106: Common Command Format – Admin and NVM Vendor Specific Commands (Optional)**

Bytes	Description
03:00	<b>Command Dword 0 (CDW0):</b> This field is common to all commands and is defined in Figure 104.
07:04	<b>Namespace Identifier (NSID):</b> This field indicates the namespace ID that this command applies to. If the namespace ID is not used for the command, then this field shall be cleared to 0h. Setting this value to FFFFFFFFh causes the command to be applied to all namespaces attached to this controller, unless otherwise specified.  The behavior of a controller in response to an inactive namespace ID for a vendor specific command is vendor specific. Specifying an invalid namespace ID in a command that uses the namespace ID shall cause the controller to abort the command with status Invalid Namespace or Format, unless otherwise specified.
15:08	Reserved
39:16	Refer to Figure 105 for the definition of these fields.
43:40	<b>Number of Dwords in Data Transfer (NDT):</b> This field indicates the number of dwords in the data transfer.
47:44	<b>Number of Dwords in Metadata Transfer (NDM):</b> This field indicates the number of dwords in the metadata transfer.
51:48	<b>Command Dword 12 (CDW12):</b> This field is command specific Dword 12.
55:52	<b>Command Dword 13 (CDW13):</b> This field is command specific Dword 13.
59:56	<b>Command Dword 14 (CDW14):</b> This field is command specific Dword 14.
63:60	<b>Command Dword 15 (CDW15):</b> This field is command specific Dword 15.

Modify section 4.6 as follows:

## 4.6 Common Completion Queue Entry

The Common Completion Queue Entry is defined in this section. An entry in the Common Completion Queue is at least 16 bytes in size. Figure 121 describes the layout of the first 16 bytes of the Common Completion Queue Entry data structure. The contents of Dword 0 and Dword 1 are command specific. If a command uses Dword 0 or Dword 1, then the definition of this these dwords is contained within the associated command definition. If a command does not use Dword 0 or Dword 1, then the unused field(s) are reserved. ~~Dword 1 is reserved.~~ Dword 2 is defined in Figure 122 and Dword 3 is defined in Figure 123. ~~Any additional I/O Command Set defined in the future may use an alternate Completion Queue entry size or format.~~

If a Completion Queue Entry is constructed via multiple writes, the Phase Tag bit shall be updated in the last write of that Completion Queue Entry.



**Figure 121: Common Completion Queue Entry Layout – Admin and ~~NVM~~ All I/O Command Sets**

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

**Figure 122: Completion Queue Entry: DW 2**

Bits	Description
31:16	<p><b>SQ Identifier (SQID):</b> Indicates the Submission Queue to which the associated command was issued. This field is used by host software when more than one Submission Queue shares a single Completion Queue to uniquely determine the command completed in combination with the Command Identifier (CID).</p> <p>This is a reserved field in NVMe over Fabrics implementations.</p>
15:00	<p><b>SQ Head Pointer (SQHD):</b> Indicates the current Submission Queue Head pointer for the Submission Queue indicated in the SQ Identifier field. This is used to indicate to the host the Submission Queue entries that have been consumed and may be re-used for new entries.</p> <p>Note: The value returned is the value of the SQ Head pointer when the completion queue entry was created. By the time host software consumes the completion queue entry, the controller may have an SQ Head pointer that has advanced beyond the value indicated.</p>

**Figure 123: Completion Queue Entry: DW 3**

Bits	Description
31:17	<p><b>Status Field (SF):</b> Indicates status for the command that is being completed. Refer to section 4.6.1.</p>
16	<p><b>Phase Tag (P):</b> Identifies whether a Completion Queue entry is new. The Phase Tag values for all Completion Queue entries shall be initialized to '0' by host software prior to setting CC.EN to '1'. When the controller places an entry in the Completion Queue, the controller shall invert the Phase Tag to enable host software to discriminate a new entry. Specifically, for the first set of completion queue entries after CC.EN is set to '1' all Phase Tags are set to '1' when they are posted. For the second set of completion queue entries, when the controller has wrapped around to the top of the Completion Queue, all Phase Tags are cleared to '0' when they are posted. The value of the Phase Tag is inverted each pass through the Completion Queue.</p> <p>This is a reserved bit in NVMe over Fabrics implementations.</p>
15:00	<p><b>Command Identifier (CID):</b> Indicates the identifier of the command that is being completed. This identifier is assigned by host software when the command is submitted to the Submission Queue. The combination of the SQ Identifier and Command Identifier uniquely identifies the command that is being completed. The maximum number of requests outstanding for a Submission Queue at one time is 64 Ki.</p>

## Modify section 4.6.1.2.1 :

### 4.6.1.2.1 Generic Command Status Definition

Completion queue entries with a Status Code type of Generic Command Status indicate a status value associated with the command that is generic across many different types of commands.

**Figure 126: Status Code – Generic Command Status Values**

Value	Description	Command Set specific	I/O Command Set(s) <sup>1</sup>
00h	<b>Successful Completion:</b> The command completed without error.	No	
01h	<b>Invalid Command Opcode:</b> A reserved coded value or an unsupported value in the command opcode field.	No	
02h	<b>Invalid Field in Command:</b> A reserved coded value or an unsupported value in a defined field (other than the opcode field). This status code should be used unless another status code is explicitly specified for a particular condition. The field may be in the command parameters as part of the Submission Queue Entry or in data structures pointed to by the command parameters.	No	
03h	<b>Command ID Conflict:</b> The command identifier is already in use. Note: It is implementation specific how many commands are searched for a conflict.	No	
04h	<b>Data Transfer Error:</b> Transferring the data or metadata associated with a command had an error.	No	
05h	<b>Commands Aborted due to Power Loss Notification:</b> Indicates that the command was aborted due to a power loss notification.	No	
06h	<b>Internal Error:</b> The command was not completed successfully due to an internal error. Details on the internal device error should be reported as an asynchronous event. Refer to Figure 145 for Internal Error Asynchronous Event Information.	No	
07h	<b>Command Abort Requested:</b> The command was aborted due to an Abort command being received that specified the Submission Queue Identifier and Command Identifier of this command (refer to section 5.1).	No	
08h	<b>Command Aborted due to SQ Deletion:</b> The command was aborted due to a Delete I/O Submission Queue request received for the Submission Queue to which the command was submitted.	No	
09h	<b>Command Aborted due to Failed Fused Command:</b> The command was aborted due to the other command in a fused operation failing.	Yes	N
0Ah	<b>Command Aborted due to Missing Fused Command:</b> The fused command was aborted due to the adjacent submission queue entry not containing a fused command that is the other command in a supported fused operation (refer to section 6.2).	Yes	N
0Bh	<b>Invalid Namespace or Format:</b> The namespace or the format of that namespace is invalid.	No	
0Ch	<b>Command Sequence Error:</b> The command was aborted due to a protocol violation in a multi-command sequence (e.g., a violation of the Security Send and Security Receive sequencing rules in the TCG Storage Synchronous Interface Communications protocol (refer to TCG Storage Architecture Core Specification)).	No	

**Figure 126: Status Code – Generic Command Status Values**

Value	Description	Command Set specific	I/O Command Set(s) <sup>1</sup>
0Dh	<b>Invalid SGL Segment Descriptor:</b> The command includes an invalid SGL Last Segment or SGL Segment descriptor. This may occur when the SGL segment pointed to by an SGL Last Segment descriptor contains an SGL Segment descriptor or an SGL Last Segment descriptor or an SGL Segment descriptor. This may occur when an SGL Last Segment descriptor contains an invalid length (i.e., a length of 0h or 1h that is not a multiple of 16).	No	
0Eh	<b>Invalid Number of SGL Descriptors:</b> There is an SGL Last Segment descriptor or an SGL Segment descriptor in a location other than the last descriptor of a segment based on the length indicated.	No	
0Fh	<b>Data SGL Length Invalid:</b> This may occur if the length of a Data SGL is too short. This may occur if the length of a Data SGL is too long and the controller does not support SGL transfers longer than the amount of data to be transferred as indicated in the SGL Support field of the Identify Controller data structure.	No	
10h	<b>Metadata SGL Length Invalid:</b> This may occur if the length of a Metadata SGL is too short. This may occur if the length of a Metadata SGL is too long and the controller does not support SGL transfers longer than the amount of data to be transferred as indicated in the SGL Support field of the Identify Controller data structure.	No	
11h	<b>SGL Descriptor Type Invalid:</b> The type of an SGL Descriptor is a type that is not supported by the controller.	No	
12h	<b>Invalid Use of Controller Memory Buffer:</b> The attempted use of the Controller Memory Buffer is not supported by the controller. Refer to section 4.7.	No	
13h	<b>PRP Offset Invalid:</b> The Offset field for a PRP entry is invalid. This may occur when there is a PRP entry with a non-zero offset after the first entry or when the Offset field in any PRP entry is not dword aligned (i.e., bits 1:0 are not cleared to 00b).	No	
14h	<b>Atomic Write Unit Exceeded:</b> The length specified exceeds the atomic write unit size.	Yes	N
15h	<b>Operation Denied:</b> The command was denied due to lack of access rights. Refer to the appropriate security specification (e.g., TCG Storage Interface Interactions specification). For media access commands, the Access Denied status code should be used instead.	No	
16h	<b>SGL Offset Invalid:</b> The offset specified in a descriptor is invalid. This may occur when using capsules for data transfers in NVMe over Fabrics implementations and an invalid offset in the capsule is specified.	No	
17h	Reserved		
18h	<b>Host Identifier Inconsistent Format:</b> The NVM subsystem detected the simultaneous use of 64-bit and 128-bit Host Identifier values on different controllers.	No	
19h	<b>Keep Alive Timer Expired:</b> The Keep Alive Timer expired.	No	
1Ah	<b>Keep Alive Timeout Invalid:</b> The Keep Alive Timeout value specified is invalid. This may be due to an attempt to specify a value of 0h on a transport that requires the Keep Alive feature to be enabled. This may be due to the value specified being too large for the associated NVMe Transport as defined in the NVMe Transport binding specification.	No	

**Figure 126: Status Code – Generic Command Status Values**

Value	Description	Command Set specific	I/O Command Set(s) <sup>1</sup>
1Bh	<b>Command Aborted due to Preempt and Abort:</b> The command was aborted due to a Reservation Acquire command with the Reservation Acquire Action (RACQA) set to 010b (Preempt and Abort).	No	
1Ch	<b>Sanitize Failed:</b> The most recent sanitize operation failed and no recovery action has been successfully completed.	No	
1Dh	<b>Sanitize In Progress:</b> The requested function (e.g., command) is prohibited while a sanitize operation is in progress. Refer to section 8.15.1.	No	
1Eh	<b>SGL Data Block Granularity Invalid:</b> The Address alignment or Length granularity for an SGL Data Block descriptor is invalid. This may occur when a controller supports dword granularity only and the lower two bits of the Address or Length are not cleared to 00b.  Note: An implementation compliant to revision 1.2.1 or earlier may use the status code value of 15h to indicate SGL Data Block Granularity Invalid.	Yes	N
1Fh	<b>Command Not Supported for Queue in CMB:</b> The implementation does not support submission of the command to a Submission Queue in the Controller Memory Buffer or command completion to a Completion Queue in the Controller Memory Buffer.  Note: Revision 1.3 and later of this specification use this status code only for Sanitize commands.	No	
20h	<b>Namespace is Write Protected:</b> The command is prohibited while the namespace is write protected as a result of a change in the namespace write protection state as defined by the Namespace Write Protection State Machine (refer to Figure 489).	No	
21h	<b>Command Interrupted:</b> Command processing was interrupted and the controller is unable to successfully complete the command. The host should retry the command.  If this status code is returned, then the controller shall clear the Do Not Retry bit to '0' in the Status field of the CQE (refer to Figure 124). The controller shall not return this status code unless the host has set the Advanced Command Retry Enable (ACRE) field to 1h in the Host Behavior Support feature (refer to section 5.21.1.22).	No	
22h	<b>Transient Transport Error:</b> A transient transport error was detected. If the command is retried on the same controller, the command is likely to succeed. A command that fails with a transient transport error four or more times should be treated as a persistent transport error that is not likely to succeed if retried on the same controller.	No	
23h to 7Fh	Reserved		
80h	<b>LBA Out of Range:</b> The command references an LBA that exceeds the size of the namespace.	Yes	N
81h	<b>Capacity Exceeded:</b> Execution of the command has caused the capacity of the namespace to be exceeded. This error occurs when the Namespace Utilization exceeds the Namespace Capacity, as reported in Figure 245.	No	
82h	<b>Namespace Not Ready:</b> The namespace is not ready to be accessed as a result of a condition other than a condition	No	

**Figure 126: Status Code – Generic Command Status Values**

Value	Description	Command Set specific	I/O Command Set(s) <sup>1</sup>
	that is reported as an Asymmetric Namespace Access condition. The Do Not Retry bit indicates whether re-issuing the command at a later time may succeed.		
83h	<b>Reservation Conflict:</b> The command was aborted due to a conflict with a reservation held on the accessed namespace. Refer to section 8.8.	No	
84h	<b>Format In Progress:</b> A Format NVM command is in progress on the namespace. The Do Not Retry bit shall be cleared to '0' to indicate that the command may succeed if resubmitted.	Yes	N
805h to BFh	<del>I/O Command Set Specific</del> Reserved		
C0h to FFh	Vendor Specific		
Key: N – NVM Command Set 1. This column is blank unless the value is I/O Command Set Specific			

**Figure 127: Status Code – Generic Command Status Values, NVM Command Set**

Value	Description
80h	<del>LBA Out of Range:</del> The command references an LBA that exceeds the size of the namespace.
81h	<del>Capacity Exceeded:</del> Execution of the command has caused the capacity of the namespace to be exceeded. This error occurs when the Namespace Utilization exceeds the Namespace Capacity, as reported in Figure 245.
82h	<del>Namespace Not Ready:</del> The namespace is not ready to be accessed as a result of a condition other than a condition that is reported as an Asymmetric Namespace Access condition. The Do Not Retry bit indicates whether re-issuing the command at a later time may succeed.
83h	<del>Reservation Conflict:</del> The command was aborted due to a conflict with a reservation held on the accessed namespace. Refer to section 8.8.
84h	<del>Format In Progress:</del> A Format NVM command is in progress on the namespace. The Do Not Retry bit shall be cleared to '0' to indicate that the command may succeed if resubmitted.
85h to BFh	Reserved

## Modify Figure 128: Status Code – Command Specific Values:

**Figure 128: Status Code – Command Specific Status Values**

Value	Description	Commands Affected
00h	Completion Queue Invalid	Create I/O Submission Queue
01h	Invalid Queue Identifier	Create I/O Submission Queue, Create I/O Completion Queue, Delete I/O Completion Queue, Delete I/O Submission Queue
02h	Invalid Queue Size	Create I/O Submission Queue, Create I/O Completion Queue
...	...	...
21h	Invalid Number of Controller Resources	Virtualization Management
22h	Invalid Resource Identifier	Virtualization Management

**Figure 128: Status Code – Command Specific Status Values**

Value	Description	Commands Affected
23h	Sanitize Prohibited While Persistent Memory Region is Enabled	Sanitize
24h	ANA Group Identifier Invalid	Namespace Management
25h	ANA Attach Failed	Namespace Attachment
...	...	...
29h	I/O Command Set Not Supported	Namespace Attachment, Namespace Management
2Ah	I/O Command Set Not Enabled	Namespace Attachment
2Bh	I/O Command Set Combination Rejected	Set Features
2Ch	Invalid I/O Command Set	Identify, Namespace Management
2Dh to 6Fh	Reserved	
70h to 7Fh	Directive Specific	NOTE 1
80h to BFh	I/O Command Set Specific	NOTE 2
C0h to FFh	Vendor Specific	
NOTES: 1. The Directives Specific range defines Directives specific status values. Refer to section 9. 2. The I/O Command Set Specific range in the NVMe over Fabrics specification defines Fabrics command specific status values.		

**Figure 129: Status Code – Command Specific Status Values, NVM Command Set**

Value	Description	Commands Affected
80h	Conflicting Attributes	Dataset Management, Read, Write
81h	Invalid Protection Information	Compare, Read, Verify, Write, Write Zeroes
82h	Attempted Write to Read Only Range	Dataset Management, Write, Write Uncorrectable, Write Zeroes
83h to BFh	Reserved	

**Figure 130: Status Code – Media and Data Integrity Error Values**

Value	Description	Command Set specific	Command Set(s)
00h to 7Fh	Reserved		
80h	<b>Write Fault:</b> The write data could not be committed to the media.	No	
81h	<b>Unrecovered Read Error:</b> The read data could not be recovered from the media.	No	
82h	<b>End-to-end Guard Check Error:</b> The command was aborted due to an end-to-end guard check failure.	No	
83h	<b>End-to-end Application Tag Check Error:</b> The command was aborted due to an end-to-end application tag check failure.	No	
84h	<b>End-to-end Reference Tag Check Error:</b> The command was aborted due to an end-to-end reference tag check failure.	No	
85h	<b>Compare Failure:</b> The command failed due to a miscompare during a Compare command.	Yes	N
86h	<b>Access Denied:</b> Access to the namespace and/or LBA range is denied due to lack of access rights. Refer to the appropriate security specification (e.g., TCG Storage Interface Interactions Specification).	No	

**Figure 130: Status Code – Media and Data Integrity Error Values**

Value	Description	Command Set specific	Command Set(s)
87h	<b>Deallocated or Unwritten Logical Block:</b> The command failed due to an attempt to read from or verify an LBA range containing a deallocated or unwritten logical block.	Yes	N
88h to BFh	<del>I/O Command Set Specific</del> Reserved		
C0h to FFh	Vendor Specific		
Key: N – NVM Command Set			

**Figure 1131: Status Code—Media and Data Integrity Error Values, NVM Command Set**

Value	Description
80h	<del>Write Fault:</del> The write data could not be committed to the media.
81h	<del>Unrecovered Read Error:</del> The read data could not be recovered from the media.
82h	<del>End-to-end Guard Check Error:</del> The command was aborted due to an end-to-end guard check failure.
83h	<del>End-to-end Application Tag Check Error:</del> The command was aborted due to an end-to-end application tag check failure.
84h	<del>End-to-end Reference Tag Check Error:</del> The command was aborted due to an end-to-end reference tag check failure.
85h	<del>Compare Failure:</del> The command failed due to a miscompare during a Compare command.
86h	<del>Access Denied:</del> Access to the namespace and/or LBA range is denied due to lack of access rights. Refer to the appropriate security specification (e.g., TCG Storage Interface Interactions specification).
87h	<del>Deallocated or Unwritten Logical Block:</del> The command failed due to an attempt to read from or verify an LBA range containing a deallocated or unwritten logical block.
88h to BFh	Reserved

**Modify Figure 139: Opcodes for Admin Commands and Figure 140: Opcodes for Admin Commands – NVM Command Set Specific as shown below:**

Figure 139 defines Admin commands ~~while Figure 140 defines I/O Command Set Specific Admin commands that are specific to the NVM Command Set (i.e., NVM Command Set Specific Admin commands).~~ Refer to Section 7.1 for mandatory, optional, and prohibited commands for the various controller types.

Figure 139: Opcodes for Admin Commands

Opcode by Field			Combined Opcode <sup>1</sup>	Namespace Identifier Used <sup>2</sup>	Command	Command Set specific	I/O Command Set <sup>8</sup>
(07)	(06:02)	(01:00)					
Generic Command	Function	Data Transfer <sup>3</sup>					
0b	000 00b	00b	00h	No	Delete I/O Submission Queue	<u>No</u>	
0b	000 00b	01b	01h	No	Create I/O Submission Queue	<u>No</u>	
0b	000 00b	10b	02h	Yes	Get Log Page	<u>No</u>	
0b	000 01b	00b	04h	No	Delete I/O Completion Queue	<u>No</u>	
0b	000 01b	01b	05h	No	Create I/O Completion Queue	<u>No</u>	
0b	000 01b	10b	06h	NOTE 6	Identify	<u>No</u>	
0b	000 10b	00b	08h	No	Abort	<u>No</u>	
0b	000 10b	01b	09h	Yes	Set Features	<u>No</u>	
0b	000 10b	10b	0Ah	Yes	Get Features	<u>No</u>	
0b	000 11b	00b	0Ch	No	Asynchronous Event Request	<u>No</u>	
0b	000 11b	01b	0Dh	Yes	Namespace Management	<u>No</u>	
0b	001 00b	00b	10h	No	Firmware Commit	<u>No</u>	
0b	001 00b	01b	11h	No	Firmware Image Download	<u>No</u>	
0b	001 01b	00b	14h	Yes	Device Self-test	<u>No</u>	
0b	001 01b	01b	15h	Yes <sup>4</sup>	Namespace Attachment	<u>No</u>	
0b	001 10b	00b	18h	No	Keep Alive	<u>No</u>	
0b	001 10b	01b	19h	Yes <sup>5</sup>	Directive Send	<u>No</u>	
0b	001 10b	10b	1Ah	Yes <sup>5</sup>	Directive Receive	<u>No</u>	
0b	001 11b	00b	1Ch	No	Virtualization Management	<u>No</u>	
0b	001 11b	01b	1Dh	No	NVMe-MI Send	<u>No</u>	
0b	001 11b	10b	1Eh	No	NVMe-MI Receive	<u>No</u>	
0b	111 11b	00b	7Ch	No	Doorbell Buffer Config	<u>No</u>	
0b	111 11b	11b	7Fh	Refer to the NVMe over Fabrics specification.			
<b>I/O Command Set Specific</b>							
4b	n/a	NOTE 3	80h to BFh		I/O Command Set specific		



**Figure 139: Opcodes for Admin Commands**

Opcode by Field			Combined Opcode <sup>1</sup>	Namespace Identifier Used <sup>2</sup>	Command	Command Set specific	I/O Command Set <sup>8</sup>
(07)	(06:02)	(01:00)					
Generic Command	Function	Data Transfer <sup>3</sup>					
1b	000 00b	00b	80h	Yes	<u>Format NVM</u>	<u>Yes</u>	<u>N</u>
1b	000 00b	01b	81h	NOTE 7	<u>Security Send</u>	<u>Yes</u>	<u>N</u>
1b	000 00b	10b	82h	NOTE 7	<u>Security Receive</u>	<u>Yes</u>	<u>N</u>
1b	000 01b	00b	84h	No	<u>Sanitize</u>	<u>Yes</u>	<u>N</u>
1b	000 01b	10b	86h	NOTE 4	<u>Get LBA Status</u>	<u>Yes</u>	<u>N</u>
<b>Vendor Specific</b>							
1b	n/a	NOTE 3	C0h to FFh		Vendor specific		

**Key:**  
N – NVM Command Set

**NOTES:**

- Opcodes not listed are reserved.
- A subset of commands use the Namespace Identifier (NSID) field. If the Namespace Identifier field is used, then the value FFFFFFFFh is supported in this field unless otherwise indicated in footnotes in this figure that a specific command does not support that value or supports that value only under specific conditions. When this field is not used, the field is cleared to 0h as described in Figure 105.
- Indicates the data transfer direction of the command. All options to the command shall transfer data as specified or transfer no data. All commands, including vendor specific commands, shall follow this convention: 00b = no data transfer; 01b = host to controller; 10b = controller to host; 11b = bidirectional.
- This command does not support the use of the Namespace Identifier (NSID) field set to FFFFFFFFh.
- Support for the Namespace Identifier field set to FFFFFFFFh depends on the Directive Operation (refer to section 9).
- Use of the Namespace Identifier field depends on the CNS value in the Identify Command as described in Figure 244.
- The use of the Namespace Identifier is Security Protocol specific.
- This column is blank unless the command is I/O Command Set Specific

**Figure 140: Opcodes for Admin Commands – NVM Command Set Specific**

Opcode (07)	Opcode (06:02)	Opcode (01:00)	Opcode <sup>4</sup>	Namespace Identifier Used <sup>2</sup>	Command
Generic Command	Function	Data Transfer <sup>3</sup>			
4b	000-00b	00b	80h	Yes	<u>Format NVM</u>
4b	000-00b	01b	81h	NOTE 4	<u>Security Send</u>
4b	000-00b	10b	82h	NOTE 4	<u>Security Receive</u>

**Figure 140: Opcodes for Admin Commands – NVM Command Set Specific**

Opcode (07)	Opcode (06:02)	Opcode (01:00)	Opcode <sup>4</sup>	Namespace Identifier Used <sup>2</sup>	Command
Generic Command	Function	Data Transfer <sup>3</sup>			
4b	000-01b	00b	84h	No	Sanitize
4b	000-01b	10b	86h	NOTE 5	Get LBA Status
NOTES: 1. <del>NVM Command Set Specific opcodes not listed are reserved.</del> 2. <del>A subset of commands use the Namespace Identifier (NSID) field. If the Namespace Identifier field is used, then unless otherwise specified, the value FFFFFFFFh is supported in this field. When this field is not used, the field is cleared to 0h as described in Figure 2.</del> 3. <del>Indicates the data transfer direction of the command. All options to the command shall transfer data as specified or transfer no data. All commands, including vendor specific commands, shall follow this convention: 00b = no data transfer; 01b = host to controller; 10b = controller to host; 11b = bidirectional.</del> 4. <del>The use of the Namespace Identifier is Security Protocol specific.</del> 5. <del>This command does not support the use of the Namespace Identifier (NSID) field set to FFFFFFFFh.</del>					

**Modify Section 5 Admin Command Set as shown below:**

## 5 Admin Command Set

The Admin Command Set defines the commands that may be submitted to the Admin Submission Queue.

The Submission Queue Entry (SQE) structure and the fields that are common to all Admin commands are defined in section 4.3. The Completion Queue Entry (CQE) structure and the fields that are common to all Admin commands are defined in section 4.6. The command specific fields in the SQE and CQE structures (i.e., SQE Command Dwords 10 to 15, ~~and~~ CQE Dword 0, ~~and CQE Dword 1~~) for the Admin Command Set are defined in this section.

**Modify Section 5.2 Asynchronous Event Request command as shown below:**

### 5.2 Asynchronous Event Request command

Asynchronous events are used to notify host software of status, error, and health information as these events occur. To enable asynchronous events to be reported by the controller, host software needs to submit one or more Asynchronous Event Request commands to the controller. The controller specifies an event to the host by completing an Asynchronous Event Request command. Host software should expect that the controller may not execute the command immediately; the command should be completed when there is an event to be reported.

The Asynchronous Event Request command is submitted by host software to enable the reporting of asynchronous events from the controller. This command has no timeout. The controller posts a completion queue entry for this command when there is an asynchronous event to report to the host. If Asynchronous

Event Request commands are outstanding when the controller is reset, then each of those commands is aborted and should not return a CQE.

All command specific fields are reserved.

Host software may submit multiple Asynchronous Event Request commands to reduce event reporting latency. The total number of simultaneously outstanding Asynchronous Event Request commands is limited by the Asynchronous Event Request Limit specified in the Identify Controller data structure in **Figure 247**.

Asynchronous events are grouped into event types. The event type information is indicated in Dword 0 of the completion queue entry for the Asynchronous Event Request command. When the controller posts a completion queue entry for an outstanding Asynchronous Event Request command and thus reports an asynchronous event, subsequent events of that event type are automatically masked by the controller until the host clears that event. An event is cleared by reading the log page associated with that event using the Get Log Page command (refer to section 5.14).

The following event types are defined:

- a) **Error event:** Indicates a general error that is not associated with a specific command (refer to Figure 145). To clear this event, host software reads the Error Information log (refer to section 5.14.1.1) using the Get Log Page command with the Retain Asynchronous Event bit cleared to '0';
- b) **SMART / Health Status event:** Indicates a SMART or health status event (refer to Figure 146). To clear this event, host software reads the SMART / Health Information log (refer to section 5.14.1.2) using the Get Log Page command with the Retain Asynchronous Event bit cleared to '0'. The SMART / Health conditions that trigger asynchronous events may be configured in the Asynchronous Event Configuration feature using the Set Features command (refer to section 5.21);
- c) **Notice event:** Indicates a general event (refer to Figure 147). To clear this event, host software reads the appropriate log page as described in Figure 147. The conditions that trigger asynchronous events may be configured in the Asynchronous Event Configuration feature using the Set Features command (see section 5.21.1.11). These notice events include:
  - A. Namespace Attribute Changed;
  - B. Firmware Activation Starting;
  - C. Telemetry Log Changed;
  - D. Asymmetric Namespace Access Change;
  - E. Predictable Latency Event Aggregate Log Change;
  - F. LBA Status Information Alert; and
  - G. Endurance Group Event Aggregate Log Page Change;
- d) **NVM I/O Command Set Specific events:** Events that are defined by an I/O eCommand sSet:
  - A. **Reservation Log Page Available event:** Indicates that one or more Reservation Notification log pages (refer to section 5.14.1.16.1) are available. To clear this event, host software reads the Reservation Notification log page using the Get Log Page command with the Retain Asynchronous Event bit cleared to '0';
  - B. **Sanitize Operation Completed event:** Indicates that a sanitize operation has completed (including any associated additional media modification, refer to the No-Deallocate Modifies Media After Sanitize field in Figure 247) without unexpected deallocation of all LBAs (refer to section 5.21.1.23) and status is available in the Sanitize Status log page (refer to section 5.14.1.17). To clear this event, host software reads the Sanitize Status log page using the Get Log Page command with the Retain Asynchronous Event bit cleared to '0'; and
  - C. **Sanitize Operation Completed With Unexpected Deallocation event:** Indicates that a sanitize operation has completed with unexpected deallocation of all LBAs (refer to section 5.21.1.23) and status is available in the Sanitize Status log page (refer to section 5.14.1.17). To clear this event, host software reads the Sanitize Status log page using the Get Log Page command with the Retain Asynchronous Event bit cleared to '0';
- and
- e) **Vendor Specific event:** Indicates a vendor specific event. To clear this event, host software reads the indicated vendor specific log page using the Get Log Page command with the Retain Asynchronous Event bit cleared to '0'.

The Sanitize Operation Completed With Unexpected Deallocation asynchronous event shall be supported if the controller supports the Sanitize Config feature (refer to section 5.21.1.23).

Asynchronous events are reported due to a new entry being added to a log page (e.g., Error Information log) or a status update (e.g., status in the SMART / Health log). A status change may be permanent (e.g., the media has become read only) or transient (e.g., the temperature reached or exceeded a threshold for a period of time). Host software should modify the event threshold or mask the event for transient and permanent status changes before issuing another Asynchronous Event Request command to avoid repeated reporting of asynchronous events.

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 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.

...

**Figure 147: Asynchronous Event Information – Notice**

Value	Description
00h	<p><b>Namespace Attribute Changed:</b> The Identify Namespace data structure (refer to Figure 245) for one or more namespaces, <b>any of the supported Identify I/O Command Set Specific Namespace data structures, or as-well-as</b> the Namespace List returned when the Identify command is issued with the CNS field set to 02h, have changed. Host software may use this event as an indication to read the Identify Namespace data structures for each namespace to determine what has changed.</p> <p>Alternatively, host software may request the Changed Namespace List (Log Identifier 04h) (refer to section 5.14.1.4) to determine which namespaces in this controller have changed information in the Identify Namespace data structure <b>or in any of the supported Identify I/O Command Set Specific Namespace data structures</b> since the last time the log page was read.</p> <p>A controller shall not send this event if:</p> <ul style="list-style-type: none"> <li>a) Namespace Utilization (refer to Figure 245) has changed, as this is a frequent event that does not require action by the host;</li> <li>b) the ANAGRPID field (refer to Figure 245) has changed; or</li> <li>c) capacity information (i.e., the NUSE field and the NVMCAP field) returned in the Identify Namespace data structure (refer to Figure 245) changed as a result of an ANA state change.</li> </ul> <p>A controller shall only send this event for changes to the Format Progress Indicator field when bits 6:0 of that field transition from a non-zero value to 0h, or from 0h to a non-zero value.</p>
...	

### Modify Section 5.13 Get Features command as shown below:

Figure 184 describes the Feature Identifiers whose attributes may be retrieved using the Get Features command. The definition of the attributes returned and the associated format is specified in the section indicated.

**Figure 184: Get Features – Feature Identifiers**

Description	Section Defining Format of Attributes Returned	I/O Command Set Specific
Arbitration	5.21.1.1	No
Power Management	5.21.1.2	No
LBA Range Type	5.21.1.3	Yes <sup>1</sup>
Temperature Threshold	5.21.1.4	No
Error Recovery	5.21.1.5	Yes <sup>1</sup>
Volatile Write Cache	5.21.1.6	No
Number of Queues	5.21.1.7	No
Interrupt Coalescing	5.21.1.8	No
Interrupt Vector Configuration	5.21.1.9	No
Write Atomicity	5.21.1.10	No
Asynchronous Event Configuration	5.21.1.11	No
Autonomous Power State Transition	5.21.1.12	No
Host Memory Buffer	5.21.1.13	No
Timestamp	5.21.1.14	No
Keep Alive Timer	5.21.1.15	No
Host Controlled Thermal Management	5.21.1.16	No
Non-Operational Power State Config	5.21.1.17	No
Read Recovery Level Config	5.21.1.18	No
Predictable Latency Mode Config	5.21.1.19	No
Predictable Latency Mode Window	5.21.1.20	No
LBA Status Information Report Interval	5.21.1.21	Yes <sup>1</sup>
Host Behavior Support	5.21.1.22	No
Sanitize Config	5.21.1.23	No
Endurance Group Event Configuration	5.21.1.24	No
<b>NVM Command Set Specific</b>		
Software Progress Marker	5.21.1.25	No
Host Identifier	5.21.1.26	No
Reservation Notification Mask	5.21.1.27	No
Reservation Persistence	5.21.1.28	No
Namespace Write Protection Config	5.21.1.29	No
I/O Command Set Profile	5.21.1.TBDi	No
Notes:		
1 This feature is used by the NVM Command Set		

#### Modify 5.14 Get Log Page command Figure 201 CDW 14:

**Figure 201: Get Log Page – Command Dword 14**

Bits	Description
31:24	<b>Command Set Identifier:</b> Refer to table X.1.
31:07	Reserved
06:00	<b>UUID Index:</b> Refer to Figure 535.

## Modify 5.14.1 Log Specific Information as shown below:

### 5.14.1 Log Specific Information

Figure 191 and Figure 192 define the Log pages that may be retrieved with the Get Log Page command and the scope of the information that is returned in those Log pages. Refer to section 7.1 for mandatory, optional, and prohibited Log pages for the various controller types.

Log pages that indicate a scope of NVM subsystem return information that is global to the NVM subsystem. Log pages that indicate a scope of controller return information that is specific to the controller that is processing the command. Log pages that indicate a scope of Namespace return information that is specific to the specified namespace. For log pages that indicate multiple scopes, the namespace identifier that is specified determines which information is returned. The definition of any individual field within a Log page may indicate a different scope that is specific to that individual field.

For Log Pages with a scope of NVM subsystem or controller (as shown in Figure 191 ~~and Figure 194~~), the controller should abort commands that specify namespace identifiers other than 0h or FFFFFFFFh with status Invalid Field in Command. Otherwise the rules for namespace identifier usage in Figure 105 apply.

**Figure 191: Get Log Page – Log Page Identifiers**

Log Identifier	Scope	Log Page Name	Reference Section
00h	Reserved		
01h	Controller	Error Information	5.15.1.1
02h	NVM subsystem <sup>1</sup>	SMART / Health Information	5.15.1.2
	Namespace <sup>2</sup>		
03h	NVM subsystem	Firmware Slot Information	5.15.1.3
04h	Controller	Changed Namespace List	5.15.1.4
05h	Controller	Commands Supported and Effects	5.15.1.5
06h	Controller <sup>3</sup>	Device Self-test <sup>5</sup>	5.15.1.6
	NVM subsystem <sup>4</sup>		
07h	Vendor Specific	Telemetry Host-Initiated <sup>5</sup>	5.15.1.7
08h	Vendor Specific	Telemetry Controller-Initiated <sup>5</sup>	5.15.1.8
09h	NVM subsystem	Endurance Group Information	5.15.1.9
0Ah	NVM subsystem	Predictable Latency Per NVM Set	5.15.1.10
0Bh	NVM subsystem	Predictable Latency Event Aggregate	5.15.1.11
0Ch	Controller	Asymmetric Namespace Access	5.15.1.12
0Dh	NVM subsystem	Persistent Event Log <sup>5</sup>	5.15.1.13
0Eh	Controller	LBA Status Information	5.15.1.14

**Figure 191: Get Log Page – Log Page Identifiers**

Log Identifier	Scope	Log Page Name	Reference Section
0Fh	NVM subsystem	Endurance Group Event Aggregate	5.15.1.15
10h to 6Fh	Reserved		
70h	Discovery (refer to the NVMe over Fabrics specification)		
71h to 7Fh	Reserved for NVMe over Fabrics implementations		
80h	Controller	Reservation Notification	5.14.1.16
81h	NVM subsystem	Sanitize Status	5.14.1.17
<del>80h-82h</del> to BFh	I/O Command Set Specific		
C0h to FFh	Vendor specific <sup>5</sup>		
KEY:			
Namespace = The log page contains information about a specific namespace.			
Controller = The log page contains information about the controller that is processing the command.			
NVM subsystem = The log page contains information about the NVM subsystem.			
Vendor Specific = The log page contains information that is vendor specific.			
NOTES:			
1. For namespace identifiers of 0h or FFFFFFFFh.			
2. For namespace identifiers other than 0h or FFFFFFFFh.			
3. Bit 0 is cleared to '0' in the DSTO field in the Identify Controller data structure (refer to Figure 247).			
4. Bit 0 is set to '1' in the DSTO field in the Identify Controller data structure.			
5. Selection of a UUID may be supported. Refer to section 8.24.			

**Figure 192: Get Log Page – Log Page Identifiers, NVM Command Set Specific**

Log Identifier	Scope	Description	Reference Section
80h	Controller	Reservation Notification	5.14.1.16.1
81h	NVM subsystem	Sanitize Status	5.14.1.16.2
82h to BFh	Reserved		
<p>KEY:</p> <p>Controller = The log page contains information about the controller that is processing the command.</p> <p>NVM subsystem = The log page contains information about the NVM subsystem.</p>			



Modify figure 194 Get Log Page – SMART / Health Information Log as shown below:

Figure 194: Get Log Page – SMART / Health Information Log

Bytes	Description																
00	<p><b>Critical Warning:</b> This field indicates critical warnings for the state of the controller. Each bit corresponds to a critical warning type; multiple bits may be set to '1'. If a bit is cleared to '0', then that critical warning does not apply. Critical warnings may result in an asynchronous event notification to the host. Bits in this field represent the current associated state and are not persistent.</p> <table> <tr> <th>Bits</th><th>Definition</th></tr> <tr> <td>7:6</td><td>Reserved</td></tr> <tr> <td>5</td><td>If set to '1', then the Persistent Memory Region has become read-only or unreliable (refer to section 4.8).</td></tr> <tr> <td>4</td><td>If set to '1', then the volatile memory backup device has failed. This field is only valid if the controller has a volatile memory backup solution.</td></tr> <tr> <td>3</td><td>If set to '1', then the media has been placed in read only mode. The controller shall not set this bit to '1' if the read-only condition on the media is a result of a change in the write protection state of a namespace (refer to section 8.19.1)).</td></tr> <tr> <td>2</td><td>If set to '1', then the NVM subsystem reliability has been degraded due to significant media related errors or any internal error that degrades NVM subsystem reliability.</td></tr> <tr> <td>1</td><td> <p>If set to '1', then a temperature is:</p> <ul style="list-style-type: none"> <li>a) greater than or equal to an over temperature threshold; or</li> <li>b) less than or equal to an under temperature threshold,</li> </ul> <p>(refer to section 5.21.1.4).</p> </td></tr> <tr> <td>0</td><td>If set to '1', then the available spare capacity has fallen below the threshold.</td></tr> </table>	Bits	Definition	7:6	Reserved	5	If set to '1', then the Persistent Memory Region has become read-only or unreliable (refer to section 4.8).	4	If set to '1', then the volatile memory backup device has failed. This field is only valid if the controller has a volatile memory backup solution.	3	If set to '1', then the media has been placed in read only mode. The controller shall not set this bit to '1' if the read-only condition on the media is a result of a change in the write protection state of a namespace (refer to section 8.19.1)).	2	If set to '1', then the NVM subsystem reliability has been degraded due to significant media related errors or any internal error that degrades NVM subsystem reliability.	1	<p>If set to '1', then a temperature is:</p> <ul style="list-style-type: none"> <li>a) greater than or equal to an over temperature threshold; or</li> <li>b) less than or equal to an under temperature threshold,</li> </ul> <p>(refer to section 5.21.1.4).</p>	0	If set to '1', then the available spare capacity has fallen below the threshold.
Bits	Definition																
7:6	Reserved																
5	If set to '1', then the Persistent Memory Region has become read-only or unreliable (refer to section 4.8).																
4	If set to '1', then the volatile memory backup device has failed. This field is only valid if the controller has a volatile memory backup solution.																
3	If set to '1', then the media has been placed in read only mode. The controller shall not set this bit to '1' if the read-only condition on the media is a result of a change in the write protection state of a namespace (refer to section 8.19.1)).																
2	If set to '1', then the NVM subsystem reliability has been degraded due to significant media related errors or any internal error that degrades NVM subsystem reliability.																
1	<p>If set to '1', then a temperature is:</p> <ul style="list-style-type: none"> <li>a) greater than or equal to an over temperature threshold; or</li> <li>b) less than or equal to an under temperature threshold,</li> </ul> <p>(refer to section 5.21.1.4).</p>																
0	If set to '1', then the available spare capacity has fallen below the threshold.																
02:01	<p><b>Composite Temperature:</b> Contains a value corresponding to a temperature in degrees Kelvin that represents the current composite temperature of the controller and namespace(s) associated with that controller. The manner in which this value is computed is implementation specific and may not represent the actual temperature of any physical point in the NVM subsystem. The value of this field may be used to trigger an asynchronous event (refer to section 5.21.1.4).</p> <p>Warning and critical overheating composite temperature threshold values are reported by the WCTEMP and CCTEMP fields in the Identify Controller data structure in Figure 247.</p>																
03	<p><b>Available Spare:</b> Contains a normalized percentage (0% to 100%) of the remaining spare capacity available.</p>																
04	<p><b>Available Spare Threshold:</b> When the Available Spare falls below the threshold indicated in this field, an asynchronous event completion may occur. The value is indicated as a normalized percentage (0% to 100%). The values 101 to 255 are reserved.</p>																
05	<p><b>Percentage Used:</b> Contains a vendor specific estimate of the percentage of NVM subsystem life used based on the actual usage and the manufacturer's prediction of NVM life. A value of 100 indicates that the estimated endurance of the NVM in the NVM subsystem has been consumed, but may not indicate an NVM subsystem failure. The value is allowed to exceed 100. Percentages greater than 254 shall be represented as 255. This value shall be updated once per power-on hour (when the controller is not in a sleep state).</p> <p>Refer to the JEDEC JESD218A standard for SSD device life and endurance measurement techniques.</p>																



**Figure 194: Get Log Page – SMART / Health Information Log**

Bytes	Description												
06	<p><b>Endurance Group Critical Warning Summary:</b> This field indicates critical warnings for the state of Endurance Groups. Each bit corresponds to a critical warning type, multiple bits may be set to '1'. If a bit is cleared to '0', then that critical warning does not apply to any Endurance Group. Critical warnings may result in an asynchronous event notification to the host. Bits in this field represent the current associated state and are not persistent.</p> <p>If a bit is set to '1' in one or more Endurance Groups, then the corresponding bit shall be set to '1' in this field.</p> <table> <tr> <th>Bits</th><th>Definition</th></tr> <tr> <td>7:4</td><td>Reserved</td></tr> <tr> <td>3</td><td>If set to '1', then the namespaces in one or more Endurance Groups have been placed in read only mode not as a result of a change in the write protection state of a namespace (refer to section 8.19.1).</td></tr> <tr> <td>2</td><td>If set to '1', then the reliability of one or more Endurance Groups has been degraded due to significant media related errors or any internal error that degrades NVM subsystem reliability.</td></tr> <tr> <td>1</td><td>Reserved</td></tr> <tr> <td>0</td><td>If set to '1', then the available spare capacity of one or more Endurance Groups has fallen below the threshold.</td></tr> </table>	Bits	Definition	7:4	Reserved	3	If set to '1', then the namespaces in one or more Endurance Groups have been placed in read only mode not as a result of a change in the write protection state of a namespace (refer to section 8.19.1).	2	If set to '1', then the reliability of one or more Endurance Groups has been degraded due to significant media related errors or any internal error that degrades NVM subsystem reliability.	1	Reserved	0	If set to '1', then the available spare capacity of one or more Endurance Groups has fallen below the threshold.
Bits	Definition												
7:4	Reserved												
3	If set to '1', then the namespaces in one or more Endurance Groups have been placed in read only mode not as a result of a change in the write protection state of a namespace (refer to section 8.19.1).												
2	If set to '1', then the reliability of one or more Endurance Groups has been degraded due to significant media related errors or any internal error that degrades NVM subsystem reliability.												
1	Reserved												
0	If set to '1', then the available spare capacity of one or more Endurance Groups has fallen below the threshold.												
31:07	Reserved												
47:32	<p><b>Data Units Read:</b> Contains the number of 512 byte data units the host has read from the controller; this value does not include metadata. This value is reported in thousands (i.e., a value of 1 corresponds to 1,000 units of 512 bytes read) and is rounded up (e.g., one indicates the that number of 512 byte data units read is from 1 to 1,000, three indicates that the number of 512 byte data units read is from 2,001 to 3,000).</p> <p>For the NVM <b>eCommand sSet</b>, logical blocks read as part of Compare, Read, and Verify operations shall be included in this value. <del>See I/O Command Set specifications for additional requirements for this field.</del></p> <p>For I/O Command Sets other than the NVM Command Set, refer to the specific I/O Command Set specification for the list of commands that affect this field.</p> <p>A value of 0h in this field indicates that the number of Data Units Read is not reported.</p>												
63:48	<p><b>Data Units Written:</b> Contains the number of 512 byte data units the host has written to the controller; this value does not include metadata. This value is reported in thousands (i.e., a value of 1 corresponds to 1,000 units of 512 bytes written) and is rounded up (e.g., one indicates that the number of 512 byte data units written is from 1 to 1,000, three indicates that the number of 512 byte data units written is from 2,001 to 3,000). When the LBA size is a value other than 512 bytes, the controller shall convert the amount of data written to 512 byte units.</p> <p>For the NVM <b>eCommand sSet</b>, logical blocks written as part of Write operations shall be included in this value. Write Uncorrectable commands shall not impact this value. <del>See I/O Command Set specifications for additional requirements for this field.</del></p> <p>For I/O Command Sets other than the NVM Command Set, refer to the specific I/O Command Set specification for the list of commands that affect this field.</p> <p>A value of 0h in this field indicates that the number of Data Units Written is not reported.</p>												

**Figure 194: Get Log Page – SMART / Health Information Log**

Bytes	Description
79:64	<p><b>Host Read Commands:</b> Contains the number of read commands completed by the controller. For the NVM <b>eCommand sSet</b>, this value is the sum of the number of Compare commands and the number of Read commands. <del>See I/O Command Set specifications for additional requirements for this field.</del></p> <p>For I/O Command Sets other than the NVM Command Set, refer to the specific I/O Command Set specification for the list of commands that affect this field.</p>
95:80	<p><b>Host Write Commands:</b> Contains the number of write commands completed by the controller. For the NVM <b>eCommand sSet</b>, this is the number of Write commands. <del>See I/O Command Set specifications for additional requirements for this field.</del></p> <p>For I/O Command Sets other than the NVM Command Set, refer to the specific I/O Command Set specification for the list of commands that affect this field.</p>
111:96	<p><b>Controller Busy Time:</b> Contains the amount of time the controller is busy with I/O commands. The controller is busy when there is a command outstanding to an I/O Queue (specifically, a command was issued via an I/O Submission Queue Tail doorbell write and the corresponding completion queue entry has not been posted yet to the associated I/O Completion Queue). This value is reported in minutes.</p>
127:112	<p><b>Power Cycles:</b> Contains the number of power cycles.</p>
143:128	<p><b>Power On Hours:</b> Contains the number of power-on hours. This may not include time that the controller was powered and in a non-operational power state.</p>
159:144	<p><b>Unsafe Shutdowns:</b> Contains the number of unsafe shutdowns. This count is incremented when a Shutdown Notification (CC.SHN) is not received prior to loss of power.</p>
175:160	<p><b>Media and Data Integrity Errors:</b> Contains the number of occurrences where the controller detected an unrecovered data integrity error. Errors such as uncorrectable ECC, CRC checksum failure, or LBA tag mismatch are included in this field. Errors introduced as a result of a Write Uncorrectable command (refer to section 6.16) may or may not be included in this field.</p>
191:176	<p><b>Number of Error Information Log Entries:</b> Contains the number of Error Information log entries over the life of the controller.</p>
195:192	<p><b>Warning Composite Temperature Time:</b> Contains the amount of time in minutes that the controller is operational and the Composite Temperature is greater than or equal to the Warning Composite Temperature Threshold (WCTEMP) field and less than the Critical Composite Temperature Threshold (CCTEMP) field in the Identify Controller data structure in Figure 247.</p> <p>If the value of the WCTEMP or CCTEMP field is 0h, then this field is always cleared to 0h regardless of the Composite Temperature value.</p>
199:196	<p><b>Critical Composite Temperature Time:</b> Contains the amount of time in minutes that the controller is operational and the Composite Temperature is greater than or equal to the Critical Composite Temperature Threshold (CCTEMP) field in the Identify Controller data structure in Figure 247.</p> <p>If the value of the CCTEMP field is 0h, then this field is always cleared to 0h regardless of the Composite Temperature value.</p>
201:200	<p><b>Temperature Sensor 1:</b> Contains the current temperature reported by temperature sensor 1. This field is defined by Figure 195.</p>
203:202	<p><b>Temperature Sensor 2:</b> Contains the current temperature reported by temperature sensor 2. This field is defined by Figure 195.</p>

**Figure 194: Get Log Page – SMART / Health Information Log**

Bytes	Description
205:204	<b>Temperature Sensor 3:</b> Contains the current temperature reported by temperature sensor 3. This field is defined by Figure 195.
207:206	<b>Temperature Sensor 4:</b> Contains the current temperature reported by temperature sensor 4. This field is defined by Figure 195.
209:208	<b>Temperature Sensor 5:</b> Contains the current temperature reported by temperature sensor 5. This field is defined by Figure 195.
211:210	<b>Temperature Sensor 6:</b> Contains the current temperature reported by temperature sensor 6. This field is defined by Figure 195.
213:212	<b>Temperature Sensor 7:</b> Contains the current temperature reported by temperature sensor 7. This field is defined by Figure 195.
215:214	<b>Temperature Sensor 8:</b> Contains the current temperature reported by temperature sensor 8. This field is defined by Figure 195.
219:216	<b>Thermal Management Temperature 1 Transition Count:</b> Contains the number of times the controller transitioned to lower power active power states or performed vendor specific thermal management actions while minimizing the impact on performance in order to attempt to reduce the Composite Temperature because of the host controlled thermal management feature (refer to section 8.4.5) (i.e., the Composite Temperature rose above the Thermal Management Temperature 1). This counter shall not wrap once the value FFFFFFFFh is reached. A value of 0h, indicates that this transition has never occurred or this field is not implemented.
223:220	<b>Thermal Management Temperature 2 Transition Count:</b> Contains the number of times the controller transitioned to lower power active power states or performed vendor specific thermal management actions regardless of the impact on performance (e.g., heavy throttling) in order to attempt to reduce the Composite Temperature because of the host controlled thermal management feature (refer to section 8.4.5) (i.e., the Composite Temperature rose above the Thermal Management Temperature 2). This counter shall not wrap once the value FFFFFFFFh is reached. A value of 0h, indicates that this transition has never occurred or this field is not implemented.
227:224	<b>Total Time For Thermal Management Temperature 1:</b> Contains the number of seconds that the controller had transitioned to lower power active power states or performed vendor specific thermal management actions while minimizing the impact on performance in order to attempt to reduce the Composite Temperature because of the host controlled thermal management feature (refer to section 8.4.5). This counter shall not wrap once the value FFFFFFFFh is reached. A value of 0h, indicates that this transition has never occurred or this field is not implemented.
231:228	<b>Total Time For Thermal Management Temperature 2:</b> Contains the number of seconds that the controller had transitioned to lower power active power states or performed vendor specific thermal management actions regardless of the impact on performance (e.g., heavy throttling) in order to attempt to reduce the Composite Temperature because of the host controlled thermal management feature (refer to section 8.4.5). This counter shall not wrap once the value FFFFFFFFh is reached. A value of 0h, indicates that this transition has never occurred or this field is not implemented.
511:232	Reserved

Figure 194: Get Log Page – SMART / Health Information Log

Bytes	Description
-------	-------------

Modify 5.14.1.4 as follows:

5.14.1.4 Changed Namespace List (Log Identifier 04h)

This log page is used to describe namespaces attached to this controller that have:

- changed information in their Identify Namespace data structure or in any of their supported Identify I/O Command Set Specific Namespace data structures since the last time the log page was read;
- been added; and
- been deleted.

The log page contains a Namespace List with up to 1,024 entries. If more than 1,024 namespaces have changed attributes since the last time the log page was read, the first entry in the log page shall be set to FFFFFFFFh and the remainder of the list shall be zero filled.

Modify 5.14.1.6 as follows:

5.14.1.6 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 and one Commands Supported and Effects data structure per I/O command based on:

- ~~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).

Modify Figure 204: Get Log Page – Endurance Group Log (Log Identifier 09h) as shown below:

Figure 204: Get Log Page – Endurance Group Log (Log Identifier 09h)

Bytes	Description
00	<p><b>Critical Warning:</b> This field indicates critical warnings for the state of the Endurance Group. Each bit corresponds to a critical warning type; multiple bits may be set to '1'. If a bit is cleared to '0', then that critical warning does not apply. Critical warnings may result in an asynchronous event notification to the host. Bits in this field represent the current associated state and are not persistent.</p> <p>If a bit is set to '1' in all Endurance Groups in the NVM subsystem, then the corresponding bit shall be set to '1' in the Critical Warning field of the SMART / Health Information log page (refer to Figure 194).</p>

**Figure 204: Get Log Page – Endurance Group Log (Log Identifier 09h)**

Bytes	Description												
	<table> <tr> <th>Bits</th><th>Definition</th></tr> <tr> <td>7:4</td><td>Reserved</td></tr> <tr> <td>3</td><td>If set to '1', then all namespaces in the Endurance Group have been placed in read only mode for reasons other than a change in the write protect state of the namespace. The controller shall not set this bit to '1' if the read-only condition on the Endurance Group is a result of a change in the write protection state of all namespaces in the Endurance Group.</td></tr> <tr> <td>2</td><td>If set to '1', then the Endurance Group reliability has been degraded due to significant media related errors or any internal error that degrades NVM subsystem reliability.</td></tr> <tr> <td>1</td><td>Reserved</td></tr> <tr> <td>0</td><td>If set to '1', then the available spare capacity of the Endurance Group has fallen below the threshold.</td></tr> </table>	Bits	Definition	7:4	Reserved	3	If set to '1', then all namespaces in the Endurance Group have been placed in read only mode for reasons other than a change in the write protect state of the namespace. The controller shall not set this bit to '1' if the read-only condition on the Endurance Group is a result of a change in the write protection state of all namespaces in the Endurance Group.	2	If set to '1', then the Endurance Group reliability has been degraded due to significant media related errors or any internal error that degrades NVM subsystem reliability.	1	Reserved	0	If set to '1', then the available spare capacity of the Endurance Group has fallen below the threshold.
Bits	Definition												
7:4	Reserved												
3	If set to '1', then all namespaces in the Endurance Group have been placed in read only mode for reasons other than a change in the write protect state of the namespace. The controller shall not set this bit to '1' if the read-only condition on the Endurance Group is a result of a change in the write protection state of all namespaces in the Endurance Group.												
2	If set to '1', then the Endurance Group reliability has been degraded due to significant media related errors or any internal error that degrades NVM subsystem reliability.												
1	Reserved												
0	If set to '1', then the available spare capacity of the Endurance Group has fallen below the threshold.												
02:01	Reserved												
03	<b>Available Spare:</b> Contains a normalized percentage (0% to 100%) of the remaining spare capacity available for the Endurance Group.												
04	<b>Available Spare Threshold:</b> If the Available Spare falls below the threshold indicated in this field, an asynchronous event completion may occur. The value is indicated as a normalized percentage (0% to 100%). The values 101 to 255 are reserved.												
05	<p><b>Percentage Used:</b> Contains a vendor specific estimate of the percentage of life used for the Endurance Group based on the actual usage and the manufacturer's prediction of NVM life. A value of 100 indicates that the estimated endurance of the NVM in the Endurance Group has been consumed, but may not indicate an NVM failure. The value is allowed to exceed 100. Percentages greater than 254 shall be represented as 255. This value shall be updated once per power-on hour when the controller is not in a sleep state.</p> <p>Refer to the JEDEC JESD218A standard for SSD device life and endurance measurement techniques.</p>												
31:06	Reserved												
47:32	<p><b>Endurance Estimate:</b> This field is an estimate of the total number of data bytes that may be written to the Endurance Group over the lifetime of the Endurance Group assuming a write amplification of 1 (i.e., no increase in the number of write operations performed by the device beyond the number of write operations requested by a host). This value is reported in billions (i.e., a value of 1 corresponds to 1,000,000,000 bytes written) and is rounded up (e.g., one indicates the number of bytes written is from 1 to 1,000,000,000, three indicates the number of bytes written is from 2,000,000,001 to 3,000,000,000).</p> <p>A value of 0h indicates that the controller does not report an Endurance Estimate.</p>												
63:48	<p><b>Data Units Read:</b> Contains the total number of data bytes that have been read from the Endurance Group. This value does not include controller reads due to internal operations such as garbage collection. This value is reported in billions (i.e., a value of 1 corresponds to 1,000,000,000 bytes read) and is rounded up (e.g., one indicates the number of bytes read is from 1 to 1,000,000,000, three indicates the number of bytes read is from 2,000,000,001 to 3,000,000,000).</p> <p>A value of 0h indicates that the controller does not report the number of Data Units Read.</p>												
79:64	<b>Data Units Written:</b> Contains the total number of data bytes that have been written to the Endurance Group. This value does not include controller writes due to internal operations such as garbage collection. This value is reported in billions (i.e., a value of 1 corresponds to 1,000,000,000 bytes written) and is rounded up (e.g., one indicates the number of bytes written												

**Figure 204: Get Log Page – Endurance Group Log (Log Identifier 09h)**

Bytes	Description
	is from 1 to 1,000,000,000, three indicates the number of bytes written is from 2,000,000,001 to 3,000,000,000). A value of 0h indicates that the controller does not report the number of Data Units Written.
95:80	<b>Media Units Written:</b> Contains the total number of data bytes that have been written to the Endurance Group including both host and controller writes (e.g., garbage collection). This value is reported in billions (i.e., a value of 1 corresponds to 1,000,000,000 bytes written) and is rounded up (e.g., one indicates the number of bytes written is from 1 to 1,000,000,000, three indicates the number of bytes written is from 2,000,000,001 to 3,000,000,000). A value of 0h indicates that controller does not report the number of Media Units Written.
111:96	<b>Host Read Commands:</b> Contains the number of read commands completed by all controllers in the NVM subsystem for the Endurance Group. For the NVM <b>eCommand sSet</b> , this is the number of Compare commands and Read commands. <b>Refer to the I/O Command Set specifications for additional requirements for this field.</b>
127:112	<b>Host Write Commands:</b> Contains the number of write commands completed by all controllers in the NVM subsystem for the Endurance Group. For the NVM <b>eCommand sSet</b> , this is the number of Write commands. <b>Refer to the I/O Command Set specifications for additional requirements for this field.</b>
143:128	<b>Media and Data Integrity Errors:</b> Contains the number of occurrences where the controller detected an unrecovered data integrity error for the Endurance Group. Errors such as uncorrectable ECC, CRC checksum failure, or LBA tag mismatch are included in this field.
159:144	<b>Number of Error Information Log Entries:</b> Contains the number of Error Information log entries over the life of the controller for the Endurance Group.
511:160	Reserved

**Modify Figure 227: Set Feature Event Data Format as shown below:**

**Figure 227: Set Feature Event Data Format**

Bytes	Description										
03:00	<p><b>Set Feature Event Layout:</b> Defines the number of Command Dwords and the amount of data in the Memory Buffer from the Set Features command associated with this event.</p> <table border="1"> <thead> <tr> <th>Bits</th><th>Definition</th></tr> </thead> <tbody> <tr> <td>31:16</td><td><b>Memory Buffer Count:</b> Defines the number of bytes from the memory buffer that are logged in the Memory Buffer field. A value of 0h indicates that the Memory Buffer field does not exist.</td></tr> <tr> <td>15:04</td><td>Reserved</td></tr> <tr> <td>03</td><td><b>Logged Command Completion Dword 0:</b> If set to '1', then Dword 0 of the command completion for the Set Features command is included in the log. If cleared to '0', then Dword 0 of the command completion command for the Set Features command is not included in the log.</td></tr> <tr> <td>02:00</td><td><b>Dword Count:</b> contains the number of consecutive Dwords starting with Dword 10 from the Set Feature command that are reported in the Command Dwords field. The values 0h and 7h are reserved.</td></tr> </tbody> </table>	Bits	Definition	31:16	<b>Memory Buffer Count:</b> Defines the number of bytes from the memory buffer that are logged in the Memory Buffer field. A value of 0h indicates that the Memory Buffer field does not exist.	15:04	Reserved	03	<b>Logged Command Completion Dword 0:</b> If set to '1', then Dword 0 of the command completion for the Set Features command is included in the log. If cleared to '0', then Dword 0 of the command completion command for the Set Features command is not included in the log.	02:00	<b>Dword Count:</b> contains the number of consecutive Dwords starting with Dword 10 from the Set Feature command that are reported in the Command Dwords field. The values 0h and 7h are reserved.
Bits	Definition										
31:16	<b>Memory Buffer Count:</b> Defines the number of bytes from the memory buffer that are logged in the Memory Buffer field. A value of 0h indicates that the Memory Buffer field does not exist.										
15:04	Reserved										
03	<b>Logged Command Completion Dword 0:</b> If set to '1', then Dword 0 of the command completion for the Set Features command is included in the log. If cleared to '0', then Dword 0 of the command completion command for the Set Features command is not included in the log.										
02:00	<b>Dword Count:</b> contains the number of consecutive Dwords starting with Dword 10 from the Set Feature command that are reported in the Command Dwords field. The values 0h and 7h are reserved.										
(Dword Count * 4)+3: 4	<p><b>Command Dwords:</b> Contains a sequential list of Command Dwords from the Set Feature command starting with Command Dword 10. The number of entries in the list is specified by the Command Dword Count field. All non-reserved Command Dwords specified by the Set Feature command for the Feature Identifier shall be logged. The Command Dwords are ordered as defined by the <b>Common Command Format for the Admin Command Set and NVM Command</b> in Figure 105.</p>										
Data Buffer Count + (Dword Count * 4)+4: (Dword Count * 4)+4	<p><b>Memory Buffer:</b> Contains the data in the memory buffer for the Set Features command.</p> <p>If the Memory Buffer Count field is cleared to a value on 0h, then this field does not exist in the logged event.</p>										
Data Buffer Count + (Dword Count * 4)+8: Data Buffer Count + (Dword Count * 4)+5	<p><b>Command Completion Dword 0:</b> If the Logged Command Completion Dword 0 bit is set to '1', then this field contains the Dword 0 value from the Set Features command completion. If the Logged Command Completion Dword 0 bit is cleared to '0', then this field is not logged.</p>										

**Modify 5.14.1.16 NVM Command Set Specific Log Page Identifiers by removing 5.14.1.16 and upleveling 5.14.1.16.1 and 5.14.1.16.2:**

#### ~~5.14.1.16 — NVM Command Set Specific Log Page Identifiers~~

~~This section describes NVM Command Set Specific log pages.~~

#### ~~5.14.1.16.1~~ **5.14.1.16 Reservation Notification (Log Identifier 80h)**

**(no changes in the text of this section)**

#### ~~5.14.1.16.2~~ 5.14.1.17 Sanitize Status (Log Identifier 81h)

(no changes in the text of this section)

Add the following section 5.21.1.TBDp:

##### 5.21.1.TBDp I/O Command Set Profile (Feature Identifier 19h)

This Feature specifies the I/O Command Sets that may be used by the controller when all supported I/O Command Sets (110b) are selected in CC.CSS. This Feature shall be implemented if bit 43 is set to '1' in the CAP register. When CC.CSS is set to any value other than 110b, then this Feature has no effect and the I/O Command Sets that may be used by the controller are specified by CC.CSS. If CC.CSS is set to any value other than 110b and the controller receives a Set Features command for this Feature, then this command has no effect and returns a status code of Successful Completion.

When all supported I/O Command Sets (110b) is selected in CC.CSS, the value of this Feature specifies the index of the I/O Command Set Combination in the Identify I/O Command Set data structure that is used. Refer to section 5.15.2.TBD for more information. The Index is specified in the I/O Command Set Combination Index field of Command Dword 11 (refer to Figure TBDp1). If any namespace attached to the controller uses an I/O Command Set that is not supported by the specified I/O Command Set combination, then the controller shall abort the command with a status code of I/O Command Set Combination Rejected. Upon successful completion of a Set Features command for this Feature, the controller transitions to using the specified I/O Command Set Combination.

**Figure TBDp1 I/O Command Set Profile – Command Dword 11**

Bits	Description
31:09	Reserved
08:00	<b>I/O Command Set Combination Index (IOCSCI):</b> This field specifies the index of the I/O Command Set Combination that is to be used. This field is used for the Set Features command only and is ignored for the Get Features command for this Feature.  The controller shall abort a command that specifies an index that corresponds to an I/O Command Set Combination that has a value of 0h with a status code of I/O Command Set Combination Rejected.

If a Get Features command is submitted for this Feature, then the attributes described in Figure TBDp2 are returned in Dword 0 of the completion queue entry for that command.

**Figure TBDp2: I/O Command Set Profile – Completion Queue Entry Dword 0**

Bits	Description
31:09	Reserved
08:00	<b>I/O Command Set Combination Index (IOCSCI):</b> This field returns the index of the currently selected I/O Command Set Combination.



## Modify Section 5.21.2 Set Features Command Completion as shown below:

### 5.21.2 Command Completion

Upon completion of the Set Features command, the controller posts a completion queue entry to the Admin Completion Queue. If a status of Successful Completion is returned, the completion queue entry shall not be posted until the controller has completed setting attributes associated with the Feature. Set Features command specific status values are defined in Figure 343.

**Figure 354: Set Features – Command Specific Status Values**

Value	Description
0Dh	<b>Feature Identifier Not Saveable:</b> The Feature Identifier specified does not support a saveable value.
0Eh	<b>Feature Not Changeable:</b> The Feature Identifier specified does not support a changeable value.
0Fh	<b>Feature Not Namespace Specific:</b> The Feature Identifier specified is not namespace specific. The Feature Identifier settings apply across all namespaces.
14h	<b>Overlapping Range:</b> This error is indicated if the LBA Range Type data structure has overlapping ranges.
2Bh	<b>I/O Command Set Combination Rejected:</b> This error indicates that the controller did not accept the request to select the requested I/O Command Set Combination.

## Modify Section 5.15 Identify command as shown below:

### 5.15 Identify command

#### 5.15.1 Identify command overview

The Identify command returns a data buffer that describes information about the NVM subsystem, the controller or the namespace(s). The data structure is 4,096 bytes in size.

The Identify command uses the Data Pointer, Command Dword 10, Command Dword 11, and Command Dword 14 fields. All other command specific fields are reserved.

**Figure 240: Identify – Data Pointer**

Bits	Description
127:00	<b>Data Pointer (DPTR):</b> This field specifies the start of the data buffer. Refer to Figure 105 for the definition of this field. If using PRPs, this field shall not be a pointer to a PRP List as the data buffer may not cross more than one page boundary.

**Figure 241: Identify – Command Dword 10**

Bits	Description
31:16	<b>Controller Identifier (CNTID):</b> This field specifies the controller identifier used as part of some Identify operations. Whether the CNTID field is used for a particular Identify operation is indicated in Figure 244. If this field is not used as part of the Identify operation, then: <ul style="list-style-type: none"><li>• host software shall clear this field to 0h for backwards compatibility (0h is a valid controller identifier); and</li><li>• the controller shall ignore this field.</li></ul> Controllers that support the Namespace Management capability (refer to section 8.12) shall support this field.
15:08	Reserved
07:00	<b>Controller or Namespace Structure (CNS):</b> This field specifies the information to be returned to the host. Refer to Figure 244.

**Figure 242: Identify – Command Dword 11**

Bits	Description
31:24	<b>Command Set Identifier (CSI):</b> This field is CNS value specific. This field specifies the I/O Command Set to be used by the command for CNS values that require a Command Set Identifier. Refer to Figure 244 for Identify command CNS values that use this field. This field shall be cleared to 0h for Identify operations with CNS values that do not use this field. Values for this field are defined by Figure X1.
<del>31</del> 23:16	Reserved
15:00	<b>NVM Set Identifier (NVMSETID):</b> This field specifies the identifier of the NVM Set. This field is used for Identify operations with a CNS value of 04h. This field should be cleared to 0h for Identify operations with other CNS values.

If the controller supports selection of a UUID by the Identify command (refer to section 8.24), then Command Dword 14 is used to specify a UUID Index value (refer to Figure 243).

**Figure 243: Identify – Command Dword 14**

Bits	Description
31:07	Reserved
06:00	<b>UUID Index:</b> Refer to Figure 498.

The data structure returned is based on the Controller or Namespace Structure (CNS) field as shown in Figure 244. If there are fewer entries to return for the data structure indicated based on CNS value, then the unused portion of the list is zero filled. If a controller does not support a CNS value, then it shall abort the command with a status of Invalid Field in Command.

Note: The CNS field was specified as a one bit field in revision 1.0 and as a two bit field in revision 1.1. Host software should only issue CNS values defined in revision 1.0 to controllers compliant with revision 1.0. Host software should only issue CNS values defined in revision 1.1 to controllers compliant with revision 1.1. The results of issuing other CNS values to controllers compliant with revision 1.0 or 1.1, respectively, are indeterminate.

The Identify Controller data structure, ~~and~~ Identify Namespace data structure, ~~and the Identify I/O~~

Command Set Specific Namespace data structure include several identifiers. The format and layout of these identifiers is described in section 7.10.

**Figure 244: Identify – CNS Values**

CNS Value	O/M 1	Definition	NSID <sup>2</sup>	CNTID <sup>3</sup>	CSI <sup>3a</sup>	Reference Section
<b>Active Namespace Management</b>						
00h	M	Identify Namespace data structure for the specified NSID <del>of the specific namespace type</del> or the common namespace capabilities <del>for the NVM Command Set.</del> <sup>6</sup>	Y	N	N	5.15.2
01h	M	Identify Controller data structure for the controller processing the command. <sup>6</sup>	N	N	N	5.15.2.2
02h	M	Active Namespace ID list.	Y	N	N	5.15.2.3
03h	M	Namespace Identification Descriptor list for the specified NSID.	Y	N	N	5.15.2.4
04h	O	An NVM Set List (refer to Figure 250) is returned to the host for up to 31 NVM Sets. The list contains entries for NVM Set identifiers greater than or equal to the value specified in the NVM Set Identifier (CDW11.NVMSETID) field.	N	N	N	5.15.2.5
05h	M	Identify I/O Command Set specific Namespace data structure for the specified NSID for the I/O Command Set specified in the CSI field. <sup>6</sup>	Y	N	Y	5.15.2.TBDa
06h	M	Identify I/O Command Set Specific Controller data structure for the controller processing the command. <sup>6</sup>	N	N	Y	5.15.2.TBDb
07h	M	Active Namespace ID list associated with the specified I/O Command Set.	Y	N	Y	5.15.2.TBDc
08h to 0Fh		Reserved				
<b>Controller and Namespace Management</b>						
10h	O <sup>4</sup>	Allocated Namespace ID list.	Y	N	N	5.15.2.6
11h	O <sup>4</sup>	Identify Namespace data structure for the specified allocated NSID.	Y	N	N	5.15.2.7
12h	O <sup>4</sup>	Controller List of controllers attached to the specified NSID.	Y	Y	N	5.15.2.8
13h	O <sup>4</sup>	Controller List of controllers that exist in the NVM subsystem.	N	Y	N	5.15.2.9
14h	O <sup>5</sup>	Primary Controller Capabilities data structure for the specified primary controller.	N	Y	N	5.15.2.10
15h	O <sup>5</sup>	Secondary Controller list of controllers associated with the primary controller processing the command.	N	Y	N	5.15.2.11
16h	O	A Namespace Granularity List (refer to Figure 255) is returned to the host for up to sixteen Namespace Granularity Entries.	N	N	N <sup>7</sup>	5.15.2.12
17h	O	A UUID List (refer to Figure 257) is returned to the host.	N	N	N	5.15.2.13

Figure 244: Identify – CNS Values

CNS Value	O/M 1	Definition	NSID <sup>2</sup>	CNTID <sup>3</sup>	CSI <sup>3a</sup>	Reference Section
1Ah	O 4	I/O Command Set specific Allocated Namespace ID list	Y	N	Y	5.15.2.TBDd
1Bh	O 4	I/O Command Set specific Identify Namespace data structure.	Y	N	Y	5.15.2.TBDe
1Ch	O	I/O Command Set data structure	N	Y	N	5.15.2.TBD
1Dh to 1Fh		Reserved				
	Future Definition					
20h to FFh			Reserved			
	<div>NOTES:</div> <div>1. O/M definition: O = Optional, M = Mandatory.</div> <div>2. The NSID field is used: Y = Yes, N = No.</div> <div>3. The CDW10.CNTID field is used: Y = Yes, N = No.</div> <div>3a. The CDW11.CSI field is used: Y = Yes, N = No.</div> <div>4. Mandatory for controllers that support the Namespace Management capability (refer to section 8.12).</div> <div>5. Mandatory for controllers that support Virtualization Enhancements (refer to section 8.5).</div> <div>6. Selection of a UUID may be supported. Refer to section 8.24.</div> <div>7. Namespace Granularity only applies to namespaces that are associated with the NVM Command Set (i.e., Command Set Identifier 0h).</div>					

The Command Set Identifier values are defined in Figure X1

Figure X1: Command Set Identifiers

Command Set Identifier Value	Description	Reference Section
00h	NVM Command Set	5.15.2.TBDa.1
01h to 2Fh	Reserved	
30h to 3Fh	Vendor specific	
40h to FFh	Reserved	

Editor's Note – Assign 01h to Key Value Command Set and 02h to Zoned Namespace

## 5.15.2 Identify Data Structures

### 5.15.2.1 NVM Command Set Identify Namespace data structure (CNS 00h)

If the NSID is active, then the NVM Command Set Identify Namespace data structure (refer to Figure 245) is returned to the host for the namespace specified in the Namespace Identifier (NSID) field ~~if it is an active~~. If the specified namespace is an inactive NSID, then the controller returns a zero filled data structure. If the specified namespace is not associated with an I/O Command Set that supports this data structure, then the controller shall abort the command with a status code of Invalid I/O Command Set.

If the controller supports the Namespace Management capability (refer to section 8.12) and the NSID field is set to FFFFFFFFh, then the controller returns an Identify Namespace data structure that specifies **NVM Command Set** capabilities that are common across namespaces for this controller. If the controller does not support the Namespace Management capability and the NSID field is set to FFFFFFFFh, then the controller shall fail the command with a status code of Invalid Namespace or Format.

**Figure 245: Identify – Identify Namespace Data Structure, NVM Command Set Specific**

Bytes	O/M <sup>1</sup>	Description
07:00	M	<b>Namespace Size (NSZE):</b> This field indicates the total size of the namespace in logical blocks. A namespace of size $n$ consists of LBA 0 through $(n - 1)$ . The number of logical blocks is based on the formatted LBA size. This field is undefined prior to the namespace being formatted.
15:08	M	<b>Namespace Capacity (NCAP):</b> This field indicates the maximum number of logical blocks that may be allocated in the namespace at any point in time. The number of logical blocks is based on the formatted LBA size. This field is undefined prior to the namespace being formatted. This field is used in the case of thin provisioning and reports a value that is smaller than or equal to the Namespace Size. Spare LBAs are not reported as part of this field.  A logical block is allocated when it is written with a Write or Write Uncorrectable command. A logical block may be deallocated using the Dataset Management, Sanitize, or Write Zeroes command.
...	...	...
187:184	O	<b>LBA Format 14 Support (LBAF14):</b> This field indicates the LBA format 14 that is supported by the controller. The LBA format field is defined in Figure 246.
191:188	O	<b>LBA Format 15 Support (LBAF15):</b> This field indicates the LBA format 15 that is supported by the controller. The LBA format field is defined in Figure 246.
319:192		Reserved
4095:384	O	Vendor Specific
NOTES:		
1. O/M definition: O = Optional, M = Mandatory.		

The LBA format data structure is described in Figure 246.

**Figure 246: Identify – LBA Format Data Structure, NVM Command Set Specific**

Bits	Description
31:26	Reserved

**Figure 246: Identify – LBA Format Data Structure, NVM Command Set Specific**

Bits	Description										
25:24	<p><b>Relative Performance (RP):</b> This field indicates the relative performance of the LBA format indicated relative to other LBA formats supported by the controller. Depending on the size of the LBA and associated metadata, there may be performance implications. The performance analysis is based on better performance on a queue depth 32 with 4 KiB read workload. The meanings of the values indicated are included in the following table.</p> <table> <tr> <th>Value</th><th>Definition</th></tr> <tr> <td>00b</td><td>Best performance</td></tr> <tr> <td>01b</td><td>Better performance</td></tr> <tr> <td>10b</td><td>Good performance</td></tr> <tr> <td>11b</td><td>Degraded performance</td></tr> </table>	Value	Definition	00b	Best performance	01b	Better performance	10b	Good performance	11b	Degraded performance
Value	Definition										
00b	Best performance										
01b	Better performance										
10b	Good performance										
11b	Degraded performance										
23:16	<p><b>LBA Data Size (LBADS):</b> This field indicates the LBA data size supported. The value is reported in terms of a power of two (<math>2^n</math>). A value smaller than 9 (i.e., 512 bytes) is not supported. If the value reported is 0h, then the LBA format is not supported / used or is not currently available.</p>										
15:00	<p><b>Metadata Size (MS):</b> This field indicates the number of metadata bytes provided per LBA based on the LBA Data Size indicated. If there is no metadata supported, then this field shall be cleared to 0h.</p> <p>If metadata is supported, then the namespace may support the metadata being transferred as part of an extended data LBA or as part of a separate contiguous buffer. If end-to-end data protection is enabled, then the first eight bytes or last eight bytes of the metadata is the protection information (refer to the DPS file in the Identify Namespace data structure).</p>										

### 5.15.2.2 Identify Controller data structure (CNS 01h)

The Identify Controller data structure (refer to Figure 247) is returned to the host for ~~this controller.~~ the controller processing the command.

**Figure 247: Identify –Identify Controller Data Structure, I/O Command Set Independent**

Bytes	O/M <sup>1</sup>	Description
01:00	M	<b>PCI Vendor ID (VID):</b> Contains the company vendor identifier that is assigned by the PCI SIG. This is the same value as reported in the ID register in section 2.1.1.
03:02	M	<b>PCI Subsystem Vendor ID (SSVID):</b> Contains the company vendor identifier that is assigned by the PCI SIG for the subsystem. This is the same value as reported in the SS register in section 2.1.17.
23:04	M	<b>Serial Number (SN):</b> Contains the serial number for the NVM subsystem that is assigned by the vendor as an ASCII string. Refer to section 7.10 for unique identifier requirements. Refer to section 1.5 for ASCII string requirements.
63:24	M	<b>Model Number (MN):</b> Contains the model number for the NVM subsystem that is assigned by the vendor as an ASCII string. Refer to section 7.10 for unique identifier requirements. Refer to section 1.5 for ASCII string requirements.
71:64	M	<b>Firmware Revision (FR):</b> Contains the currently active firmware revision for the NVM subsystem. This is the same revision information that may be retrieved with the Get Log Page command, refer to section 5.14.1.3. Refer to section 1.5 for ASCII string requirements.

**Figure 247: Identify –Identify Controller Data Structure, I/O Command Set Independent**

Bytes	O/M <sup>1</sup>	Description
72	M	<b>Recommended Arbitration Burst (RAB):</b> This is the recommended Arbitration Burst size. The value is in commands and is reported as a power of two ( $2^n$ ). This is the same units as the Arbitration Burst size. Refer to section 4.13.
75:73	M	<b>IEEE OUI Identifier (IEEE):</b> Contains the Organization Unique Identifier (OUI) for the controller vendor. The OUI shall be a valid IEEE/RAC assigned identifier that may be registered at <a href="http://standards.ieee.org/develop/regauth/oui/public.html">http://standards.ieee.org/develop/regauth/oui/public.html</a> .
76	O	<p><b>Controller Multi-Path I/O and Namespace Sharing Capabilities (CMIC):</b> This field specifies multi-path I/O and namespace sharing capabilities of the controller and NVM subsystem.</p> <p>Bits 7:4 are reserved.</p> <p>Bit 3 if set to '1', then the NVM subsystem supports Asymmetric Namespace Access Reporting (refer to section 8.20). If cleared to '0', then the NVM subsystem does not support Asymmetric Namespace Access Reporting.</p> <p>Bit 2 if set to '1', then the controller is associated with an SR-IOV Virtual Function. If cleared to '0', then the controller is associated with a PCI Function or a Fabrics connection.</p> <p>Bit 1 if set to '1', then the NVM subsystem may contain two or more controllers. If cleared to '0', then the NVM subsystem contains only a single controller. As described in section 1.4.1, an NVM subsystem that contains multiple controllers may be used by multiple hosts, or may provide multiple paths for a single host.</p> <p>Bit 0 if set to '1', then the NVM subsystem may contain more than one NVM subsystem port. If cleared to '0', then the NVM subsystem contains only a single NVM subsystem port.</p>
77	M	<p><b>Maximum Data Transfer Size (MDTS):</b> This field indicates the maximum data transfer size for a command that transfers data between memory accessible by the host (e.g., host memory, Controller Memory Buffer (refer to section 4.7)) and the controller. The host should not submit a command that exceeds this maximum data transfer size. If a command is submitted that exceeds this transfer size, then the command is aborted with a status of Invalid Field in Command. The value is in units of the minimum memory page size (CAP.MPSMIN) and is reported as a power of two (<math>2^n</math>). A value of 0h indicates that there is no maximum data transfer size. This field includes the length of metadata, if metadata is interleaved with the logical block data. This field does not apply to commands that do not transfer data between memory accessible by the host and the controller (e.g., the Verify command, the Write Uncorrectable command, and the Write Zeroes command); there is no maximum data transfer size for those commands.</p> <p>If SGL Bit Bucket descriptors are supported, their lengths shall be included in determining if a command exceeds the Maximum Data Transfer Size for destination data buffers. Their length in a source data buffer is not included for a Maximum Data Transfer Size calculation.</p>
79:78	M	<b>Controller ID (CNTLID):</b> Contains the NVM subsystem unique controller identifier associated with the controller.
83:80	M	<b>Version (VER):</b> This field contains the value reported in the Version register defined in section 3.1.2. Implementations compliant to revision 1.2 or later of this specification shall report a non-zero value in this field.
87:84	M	<b>RTD3 Resume Latency (RTD3R):</b> This field indicates the expected latency in microseconds to resume from Runtime D3 (RTD3). Refer to section 8.4.4. A value of 0h indicates RTD3 Resume Latency is not reported.
91:88	M	<b>RTD3 Entry Latency (RTD3E):</b> This field indicates the typical latency in microseconds to enter Runtime D3 (RTD3). Refer to section 8.4.4. A value of 0h indicates RTD3 Entry Latency is not reported.

**Figure 247: Identify –Identify Controller Data Structure, I/O Command Set Independent**

Bytes	O/M <sup>1</sup>	Description
95:92	M	<p><b>Optional Asynchronous Events Supported (OAES):</b> This field indicates the optional asynchronous events supported by the controller. A controller shall not send optional asynchronous events before they are enabled by host software.</p> <p>Bits 31:15 are reserved.</p> <p>Bit 14 is set to '1' if the controller supports the Endurance Group Event Aggregate Log Page Change Notices event. If cleared to '0', then the controller does not support the Endurance Group Event Aggregate Log Page Change Notices event.</p> <p>Bit 13 is set to '1' if the controller supports the LBA Status Information Notices event. If cleared to '0', then the controller does not support the LBA Status Information Notices event.</p> <p>Bit 12 is set to '1' if the controller supports the Predictable Latency Event Aggregate Log Change Notices event. If cleared to '0', then the controller does not support the Predictable Latency Event Aggregate Log Change Notices event.</p> <p>Bit 11 is set to '1' if the controller supports sending Asymmetric Namespace Access Change Notices. If cleared to '0', then the controller does not support the Asymmetric Namespace Access Change Notices event.</p> <p>Bit 10 is reserved.</p> <p>Bit 9 is set to '1' if the controller supports the Firmware Activation Notices event. If cleared to '0', then the controller does not support the Firmware Activation Notices event.</p> <p>Bit 8 is set to '1' if the controller supports the Namespace Attribute Notices event and the associated Changed Namespace List log page. If cleared to '0', then the controller does not support the Namespace Attribute Notices event nor the associated Changed Namespace List log page.</p> <p>Bits 7:0 are reserved.</p>



**Figure 247: Identify –Identify Controller Data Structure, I/O Command Set Independent**

Bytes	O/M <sup>1</sup>	Description
99:96	M	<p><b>Controller Attributes (CTRATT):</b> This field indicates attributes of the controller.</p> <p>Bits 31:10 are reserved.</p> <p>Bit 9 (<b>UUID List</b>): If set to '1', then the controller supports reporting of a UUID List (refer to Figure 257). If cleared to '0', then the controller does not support reporting of a UUID List (refer to section 8.24).</p> <p>Bit 8 (<b>SQ Associations</b>): If set to '1', then the controller supports SQ Associations (refer to section 8.23). If cleared to '0', then the controller does not support SQ Associations.</p> <p>Bit 7 (<b>Namespace Granularity</b>): If set to '1', then the controller supports reporting of Namespace Granularity (refer to section 5.15.2.12). If cleared to '0', the controller does not support reporting of Namespace Granularity. If the Namespace Management capability (refer to section 8.12) is not supported, then this bit shall be cleared to '0'.</p> <p>Bit 6 (<b>Traffic Based Keep Alive Support – TBKAS</b>): If set to '1', then the controller supports restarting the Keep Alive Timer if an Admin command or an I/O command is processed during the Keep Alive Timeout Interval (refer to section 7.12.2). If cleared to '0', then the controller supports restarting the Keep Alive Timer only if a Keep Alive command is processed during the Keep Alive Timeout Interval (refer to section 7.12.1).</p> <p>Bit 5 (<b>Predictable Latency Mode</b>): If set to '1', then the controller supports Predictable Latency Mode (refer to section 8.18). If cleared to '0', then the controller does not support Predictable Latency Mode.</p> <p>Bit 4 (<b>Endurance Groups</b>): If set to '1', then the controller supports Endurance Groups (refer to section 8.17). If cleared to '0', then the controller does not support Endurance Groups.</p> <p>Bit 3 (<b>Read Recovery Levels</b>): If set to '1', then the controller supports Read Recovery Levels (refer to section 8.16). If cleared to '0', then the controller does not support Read Recovery Levels.</p> <p>Bit 2 (<b>NVM Sets</b>): If set to '1', then the controller supports NVM Sets (refer to section 4.9). If cleared to '0', then the controller does not support NVM Sets.</p> <p>Bit 1 (<b>Non-Operational Power State Permissive Mode</b>): If set to '1', then the controller supports host control of whether the controller may temporarily exceed the power of a non-operational power state for the purpose of executing controller initiated background operations in a non-operational power state (i.e., Non-Operational Power State Permissive Mode supported). If cleared to '0', then the controller does not support host control of whether the controller may exceed the power of a non-operational state for the purpose of executing controller initiated background operations in a non-operational state (i.e., Non-Operational Power State Permissive Mode not supported). Refer to section 5.21.1.17.</p> <p>Bit 0 if set to '1', then the controller supports a 128-bit Host Identifier. Bit 0 if cleared to '0', then the controller does not support a 128-bit Host Identifier.</p>

**Figure 247: Identify –Identify Controller Data Structure, I/O Command Set Independent**

Bytes	O/M <sup>1</sup>	Description																																	
101:100	O	<b>Read Recovery Levels Supported (RRLS):</b> 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.																																	
		Bit	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 <sup>1</sup>	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 <sup>1</sup>
		Bit	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 <sup>1</sup>																																
		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 <sup>1</sup>																																
		110:102		Reserved																															
111	M	<b>Controller Type (CNTRLTYPE):</b> This field specifies the controller type. A value of 0h indicates that the controller type is not reported.																																	
		Implementations compliant to version 1.4 or later of this specification shall report a controller type (i.e., the value 0h is reserved and shall not be used). Implementations compliant to an earlier specification version may report a value of 0h to indicate that a controller type is not reported.																																	
		Value	Controller Type	0h	Reserved (controller type not reported)	1h	I/O Controller	2h	Discovery Controller	3h	Administrative Controller	4h to FFh	Reserved																						
		Value	Controller Type																																
		0h	Reserved (controller type not reported)																																
		1h	I/O Controller																																
		2h	Discovery Controller																																
		3h	Administrative Controller																																
4h to FFh	Reserved																																		

**Figure 247: Identify –Identify Controller Data Structure, I/O Command Set Independent**

Bytes	O/M <sup>1</sup>	Description
127:112	O	<p><b>FRU Globally Unique Identifier (FGUID):</b> This field contains a 128-bit value that is globally unique for a given Field Replaceable Unit (FRU). Refer to the NVM Express™ Management Interface (NVMe-MI™) specification for the definition of a FRU. This field remains fixed throughout the life of the FRU. This field shall contain the same value for each controller associated with a given FRU.</p> <p>This field uses the EUI-64 based 16-byte designator format. Bytes 122:120 contain the 24-bit Organizationally Unique Identifier (OUI) value assigned by the IEEE Registration Authority. Bytes 127:123 contain an extension identifier assigned by the corresponding organization. Bytes 119:112 contain the vendor specific extension identifier assigned by the corresponding organization. Refer to the IEEE EUI-64 guidelines for more information. This field is big endian (refer to section 7.10).</p> <p>When not implemented, this field contains a value of 0h.</p>
129:128	O	<p><b>Command Retry Delay Time 1 (CRDT1):</b> If the Do Not Retry (DNR) bit is cleared to '0' in the CQE and the Command Retry Delay (CRD) field is set to 01b in the CQE, then this value indicates the command retry delay time in units of 100 milliseconds.</p>
131:130	O	<p><b>Command Retry Delay Time 2 (CRDT2):</b> If the DNR bit is cleared to '0' in the CQE and the CRD field is set to 10b in the CQE, then this value indicates the command retry delay time in units of 100 milliseconds.</p>
133:132	O	<p><b>Command Retry Delay Time 3 (CRDT3):</b> If the DNR bit is cleared to '0' in the CQE and CRD field is set to 11b in the CQE, then this value indicates the command retry delay time in units of 100 milliseconds.</p>
239:134		Reserved
255:240		Refer to the NVMe Management Interface Specification for definition.

**Figure 247: Identify –Identify Controller Data Structure, I/O Command Set Independent**

Bytes	O/M <sup>1</sup>	Description
257:256	M	<p><b>Optional Admin Command Support (OACS):</b> This field indicates the optional Admin commands and features supported by the controller. Refer to section 5.</p> <p>Bits 15:10 are reserved.</p> <p>Bit 9 if set to '1', then the controller supports the Get LBA Status capability (refer to section 8.22). If cleared to '0', then the controller does not support the Get LBA Status capability.</p> <p>Bit 8 if set to '1', then the controller supports the Doorbell Buffer Config command. If cleared to '0', then the controller does not support the Doorbell Buffer Config command.</p> <p>Bit 7 if set to '1', then the controller supports the Virtualization Management command. If cleared to '0', then the controller does not support the Virtualization Management command.</p> <p>Bit 6 if set to '1', then the controller supports the NVMe-MI Send and NVMe-MI Receive commands. If cleared to '0', then the controller does not support the NVMe-MI Send and NVMe-MI Receive commands.</p> <p>Bit 5 if set to '1', then the controller supports Directives. If cleared to '0', then the controller does not support Directives. A controller that supports Directives shall support the Directive Send and Directive Receive commands. Refer to section 9.</p> <p>Bit 4 if set to '1', then the controller supports the Device Self-test command. If cleared to '0', then the controller does not support the Device Self-test command.</p> <p>Bit 3 if set to '1', then the controller supports the Namespace Management capability (refer to section 8.12). If cleared to '0', then the controller does not support the Namespace Management capability.</p> <p>Bit 2 if set to '1', then the controller supports the Firmware Commit and Firmware Image Download commands. If cleared to '0', then the controller does not support the Firmware Commit and Firmware Image Download commands.</p> <p>Bit 1 if set to '1', then the controller supports the Format NVM command. If cleared to '0', then the controller does not support the Format NVM command.</p> <p>Bit 0 if set to '1', then the controller supports the Security Send and Security Receive commands. If cleared to '0', then the controller does not support the Security Send and Security Receive commands.</p>
258	M	<p><b>Abort Command Limit (ACL):</b> This field is used to convey the maximum number of concurrently executing Abort commands supported by the controller (refer to section 5.1). This is a 0's based value. It is recommended that implementations support concurrent execution of a minimum of four Abort commands.</p>
259	M	<p><b>Asynchronous Event Request Limit (AERL):</b> This field is used to convey the maximum number of concurrently outstanding Asynchronous Event Request commands supported by the controller (refer to section 5.2). This is a 0's based value. It is recommended that implementations support a minimum of four Asynchronous Event Request Limit commands outstanding simultaneously.</p>

**Figure 247: Identify –Identify Controller Data Structure, I/O Command Set Independent**

Bytes	O/M <sup>1</sup>	Description
260	M	<p><b>Firmware Updates (FRMW):</b> This field indicates capabilities regarding firmware updates. Refer to section 8.1 for more information on the firmware update process.</p> <p>Bits 7:5 are reserved.</p> <p>Bit 4 if set to '1' indicates that the controller supports firmware activation without a reset. If cleared to '0', then the controller requires a reset for firmware to be activated.</p> <p>Bits 3:1 indicate the number of firmware slots that the controller supports. This field shall specify a value from one to seven, indicating that at least one firmware slot is supported and up to seven maximum. This corresponds to firmware slots 1 through 7.</p> <p>Bit 0 if set to '1' indicates that the first firmware slot (i.e., slot 1) is read only. If cleared to '0', then the first firmware slot (i.e., slot 1) is read/write. Implementations may choose to have a baseline read only firmware image.</p>
261	M	<p><b>Log Page Attributes (LPA):</b> This field indicates optional attributes for log pages that are accessed via the Get Log Page command.</p> <p>Bits 7:5 are reserved.</p> <p>Bit 4 if set to '1', then the controller supports the Persistent Event log. If cleared to '0', then the controller does not support the Persistent Event log.</p> <p>Bit 3 if set to '1', then the controller supports the Telemetry Host-Initiated and Telemetry Controller-Initiated log pages and sending Telemetry Log Notices. If cleared to '0', then the controller does not support the Telemetry Host-Initiated and Telemetry Controller-Initiated log pages and Telemetry Log Notice events.</p> <p>Bit 2 if set to '1', then the controller supports extended data for the Get Log Page command (including extended Number of Dwords and Log Page Offset fields). Bit 2 if cleared to '0', then the controller does not support extended data for the Get Log Page command.</p> <p>Bit 1 if set to '1', then the controller supports the Commands Supported and Effects log page. Bit 1 if cleared to '0', then the controller does not support the Commands Supported and Effects log page.</p> <p>Bit 0 if set to '1', then the controller supports the SMART / Health Information log page on a per namespace basis. If cleared to '0', then the controller does not support the SMART / Health Information log page on a per namespace basis.</p>
262	M	<p><b>Error Log Page Entries (ELPE):</b> This field indicates the maximum number of Error Information log entries that are stored by the controller. This field is a 0's based value.</p>
263	M	<p><b>Number of Power States Support (NPSS):</b> This field indicates the number of NVM Express power states supported by the controller. This is a 0's based value. Refer to section 8.4.</p> <p>Power states are numbered sequentially starting at power state 0. A controller shall support at least one power state (i.e., power state 0) and may support up to 31 additional power states (i.e., up to 32 total).</p>
264	M	<p><b>Admin Vendor Specific Command Configuration (AVSCC):</b> This field indicates the configuration settings for Admin Vendor Specific command handling. Refer to section 8.7.</p> <p>Bits 7:1 are reserved.</p> <p>Bit 0 if set to '1' indicates that all Admin Vendor Specific Commands use the format defined in Figure 106. If cleared to '0' indicates that the format of all Admin Vendor Specific Commands are vendor specific.</p>
265	O	<p><b>Autonomous Power State Transition Attributes (APSTA):</b> This field indicates the attributes of the autonomous power state transition feature. Refer to section 8.4.2.</p> <p>Bits 7:1 are reserved.</p> <p>Bit 0 if set to '1', then the controller supports autonomous power state transitions. If cleared to '0', then the controller does not support autonomous power state transitions.</p>

**Figure 247: Identify –Identify Controller Data Structure, I/O Command Set Independent**

Bytes	O/M <sup>1</sup>	Description
267:266	M	<p><b>Warning Composite Temperature Threshold (WCTEMP):</b> This field indicates the minimum Composite Temperature field value (reported in the SMART / Health Information log in Figure 194) that indicates an overheating condition during which controller operation continues. Immediate remediation is recommended (e.g., additional cooling or workload reduction). The platform should strive to maintain a composite temperature less than this value.</p> <p>A value of 0h in this field indicates that no warning temperature threshold value is reported by the controller. Implementations compliant to revision 1.2 or later of this specification shall report a non-zero value in this field.</p> <p>It is recommended that implementations report a value of 0157h in this field.</p>
269:268	M	<p><b>Critical Composite Temperature Threshold (CCTEMP):</b> This field indicates the minimum Composite Temperature field value (reported in the SMART / Health Information log in Figure 194) that indicates a critical overheating condition (e.g., may prevent continued normal operation, possibility of data loss, automatic device shutdown, extreme performance throttling, or permanent damage).</p> <p>A value of 0h in this field indicates that no critical temperature threshold value is reported by the controller. Implementations compliant to revision 1.2 or later of this specification shall report a non-zero value in this field.</p>
271:270	O	<p><b>Maximum Time for Firmware Activation (MTFA):</b> Indicates the maximum time the controller temporarily stops processing commands to activate the firmware image. This field shall be valid if the controller supports firmware activation without a reset. This field is specified in 100 millisecond units. A value of 0h indicates that the maximum time is undefined.</p>
275:272	O	<p><b>Host Memory Buffer Preferred Size (HMPRE):</b> This field indicates the preferred size that the host is requested to allocate for the Host Memory Buffer feature in 4 KiB units. This value shall be greater than or equal to the Host Memory Buffer Minimum Size. If this field is non-zero, then the Host Memory Buffer feature is supported. If this field is cleared to 0h, then the Host Memory Buffer feature is not supported.</p>
279:276	O	<p><b>Host Memory Buffer Minimum Size (HMMIN):</b> This field indicates the minimum size that the host is requested to allocate for the Host Memory Buffer feature in 4 KiB units. If this field is cleared to 0h, then the host is requested to allocate any amount of host memory possible up to the HMPRE value.</p>
295:280	O	<p><b>Total NVM Capacity (TNVMCAP):</b> This field indicates the total NVM capacity in the NVM subsystem. The value is in bytes. This field shall be supported if the Namespace Management capability (refer to section 8.12) is supported.</p>
311:296	O	<p><b>Unallocated NVM Capacity (UNVMCAP):</b> This field indicates the unallocated NVM capacity in the NVM subsystem. The value is in bytes. This field shall be supported if the Namespace Management capability (refer to section 8.12) is supported.</p>

**Figure 247: Identify –Identify Controller Data Structure, I/O Command Set Independent**

Bytes	O/M <sup>1</sup>	Description																	
315:312	O	<b>Replay Protected Memory Block Support (RPMBS):</b> This field indicates if the controller supports one or more Replay Protected Memory Blocks (RPMBS) and the capabilities. Refer to section 8.10.																	
		Bits	Description	31:24	<b>Access Size:</b> 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 this 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.	23:16	<b>Total Size:</b> 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.	15:06	Reserved	05:03	<b>Authentication Method:</b> 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	<b>Number of RPMB Units:</b> 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	<b>Access Size:</b> 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 this 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.																
		23:16	<b>Total Size:</b> 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.																
		15:06	Reserved																
05:03	<b>Authentication Method:</b> 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	<b>Number of RPMB Units:</b> 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.																		
317:316	O	<b>Extended Device Self-test Time (EDSTT):</b> If the Device Self-test command is supported, then this field indicates the nominal amount of time in one minute units that the controller takes to complete an extended device self-test operation when in power state 0. If the Device Self-test command is not supported, then this field is reserved.																	
318	O	<b>Device Self-test Options (DSTO):</b> This field indicates the optional Device Self-test command or operation behaviors supported by the controller or NVM subsystem.  Bits 7:1 are reserved.  Bit 0 if set to '1', then the NVM subsystem supports only one device self-test operation in progress at a time. If cleared to '0', then the NVM subsystem supports one device self-test operation per controller at a time.																	
319	M	<b>Firmware Update Granularity (FWUG):</b> This field indicates the granularity and alignment requirement of the firmware image being updated by the Firmware Image Download command (refer to section 5.12). If the values specified in the NUMD field or the OFST field in the Firmware Image Download command do not conform to this granularity and alignment requirement, then the firmware update may fail with status of Invalid Field in Command. For the broadest interoperability with host software, it is recommended that the controller set this value to the lowest value possible.  The value is reported in 4 KiB units (e.g., 1h corresponds to 4 KiB, 2h corresponds to 8 KiB). A value of 0h indicates that no information on granularity is provided. A value of FFh indicates there is no restriction (i.e., any granularity and alignment in dwords is allowed).																	

**Figure 247: Identify –Identify Controller Data Structure, I/O Command Set Independent**

Bytes	O/M <sup>1</sup>	Description
321:320	M	<b>Keep Alive Support (KAS):</b> This field indicates the granularity of the Keep Alive Timer in 100 millisecond units (refer to section 7.12). If this field is cleared to 0h, then the Keep Alive feature is not supported. The Keep Alive feature shall be supported for NVMe over Fabrics implementations as described in section 7.12.
323:322	O	<b>Host Controlled Thermal Management Attributes (HCTMA):</b> This field indicates the attributes of the host controlled thermal management feature. Refer to section 8.4.5.  Bits 15:1 are reserved.  Bit 0 if set to '1', then the controller supports host controlled thermal management. If cleared to '0', then the controller does not support host controlled thermal management. If this bit is set to '1', then the controller shall support the Set Features command and Get Features command with the Feature Identifier field set to 10h.
325:324	O	<b>Minimum Thermal Management Temperature (MNTMT):</b> This field indicates the minimum temperature, in degrees Kelvin, that the host may request in the Thermal Management Temperature 1 field and Thermal Management Temperature 2 field of a Set Features command with the Feature Identifier field set to 10h. A value of 0h indicates that the controller does not report this field or the host controlled thermal management feature (refer to section 8.4.5) is not supported.
327:326	O	<b>Maximum Thermal Management Temperature (MXTMT):</b> This field indicates the maximum temperature, in degrees Kelvin, that the host may request in the Thermal Management Temperature 1 field and Thermal Management Temperature 2 field of the Set Features command with the Feature Identifier set to 10h. A value of 0h indicates that the controller does not report this field or the host controlled thermal management feature is not supported.



331:328	O	<b>Sanitize Capabilities (SANICAP):</b> This field indicates attributes for sanitize operations. If the Sanitize command is supported, then this field shall be non-zero. If the Sanitize command is not supported, then this field shall be cleared to 0h. Refer to section 8.15.											
		Bits	Description										
		31:30	<p><b>No-Deallocate Modifies Media After Sanitize (NODMMAS):</b> This field indicates if media is additionally modified by the NVMe controller after a sanitize operation successfully completes that had been started a Sanitize command with the No-Deallocate After Sanitize bit set to '1'.</p> <p>The work required for the associated additional media modification is included both in the estimated time for each sanitize operation and in the Sanitize Progress field (refer to Figure 238).</p> <table><tr><th>Value</th><th>Definition</th></tr><tr><td>00b</td><td>Additional media modification after sanitize operation completes successfully is not defined. Only controllers compliant with versions 1.3 and earlier of the specification shall be allowed to return this value.</td></tr><tr><td>01b</td><td>Media is not additionally modified by the NVMe controller after sanitize operation completes successfully.</td></tr><tr><td>10b</td><td>Media is additionally modified by the NVMe controller after sanitize operation completes successfully. The Sanitize Operation Completed event does not occur until the additional media modification associated with this field has completed.</td></tr><tr><td>11b</td><td>Reserved</td></tr></table>	Value	Definition	00b	Additional media modification after sanitize operation completes successfully is not defined. Only controllers compliant with versions 1.3 and earlier of the specification shall be allowed to return this value.	01b	Media is not additionally modified by the NVMe controller after sanitize operation completes successfully.	10b	Media is additionally modified by the NVMe controller after sanitize operation completes successfully. The Sanitize Operation Completed event does not occur until the additional media modification associated with this field has completed.	11b	Reserved
		Value	Definition										
		00b	Additional media modification after sanitize operation completes successfully is not defined. Only controllers compliant with versions 1.3 and earlier of the specification shall be allowed to return this value.										
01b	Media is not additionally modified by the NVMe controller after sanitize operation completes successfully.												
10b	Media is additionally modified by the NVMe controller after sanitize operation completes successfully. The Sanitize Operation Completed event does not occur until the additional media modification associated with this field has completed.												
11b	Reserved												
If bits 2:0 of the SANICAP field are cleared to 000b, then the controller shall clear this field to 00b.													
29	<p><b>No-Deallocate Inhibited (NDI):</b> If set to '1' and the No-Deallocate Response Mode bit is set to '1', then the controller deallocates after the sanitize operation even if the No-Deallocate After Sanitize bit is set to '1' in a Sanitize command.</p> <p>If:</p> <ul style="list-style-type: none"><li>a) set to '1';</li><li>b) the No-Deallocate Response Mode bit (refer to Figure 314) is cleared to '0'; and</li><li>c) the No-Deallocate After Sanitize bit is set to '1' in a Sanitize command,</li></ul> <p>then the controller aborts the Sanitize command with a status of Invalid Field in Command.</p> <p>If cleared to '0', then the controller supports the No-Deallocate After Sanitize bit in a Sanitize command.</p> <p>If bits 2:0 of the SANICAP field are cleared to 0h, then the controller shall clear this bit to '0'.</p>												
28:03	Reserved												
02	<p><b>Overwrite Support (OWS):</b> If set to '1', then the controller supports the Overwrite sanitize operation. If cleared to '0',</p>												

**Figure 247: Identify –Identify Controller Data Structure, I/O Command Set Independent**

Bytes	O/M <sup>1</sup>	Description
		<div>then the controller does not support the Overwrite sanitize operation.</div> <div>01 <b>Block Erase Support (BES):</b> If set to '1', then the controller supports the Block Erase sanitize operation. If cleared to '0', then the controller does not support the Block Erase sanitize operation.</div> <div>00 <b>Crypto Erase Support (CES):</b> If set to '1', then the controller supports the Crypto Erase sanitize operation. If cleared to '0', then the controller does not support the Crypto Erase sanitize operation.</div>
335:332	O	<b>Host Memory Buffer Minimum Descriptor Entry Size (HMMINDS):</b> This field indicates the minimum usable size of a Host Memory Buffer Descriptor Entry in 4 KiB units. If this field is cleared to 0h, then the controller does not indicate any limitations on the Host Memory Buffer Descriptor Entry size.
337:336	O	<b>Host Memory Maximum Descriptors Entries (HMMAXD):</b> This field indicates the number of usable Host Memory Buffer Descriptor Entries. If this field is cleared to 0h, then the controller does not indicate a maximum number of Host Memory Buffer Descriptor Entries.
339:338	O	<b>NVM Set Identifier Maximum (NSETIDMAX):</b> This field defines the maximum value of a valid NVM Set Identifier for any controller in the NVM subsystem. The number of NVM Sets supported by the NVM subsystem is less than or equal to NSETIDMAX.
341:340	O	<b>Endurance Group Identifier Maximum (ENDGIDMAX):</b> This field defines the maximum value of a valid Endurance Group Identifier for any controller in the NVM subsystem. The number of Endurance Groups supported by the NVM subsystem is less than or equal to ENDGIDMAX.
342	O	<b>ANA Transition Time (ANATT):</b> This field indicates the maximum amount of time, in seconds, for a transition between ANA states or the maximum amount of time, in seconds, that the controller reports the ANA change state. If the controller supports Asymmetric Namespace Access Reporting (refer to the CMIC field), then this field shall be set to a non-zero value. If the controller does not support Asymmetric Namespace Access Reporting, then this field shall be cleared to 0h. Refer to section 8.21.4.

**Figure 247: Identify –Identify Controller Data Structure, I/O Command Set Independent**

Bytes	O/M <sup>1</sup>	Description
343	O	<p><b>Asymmetric Namespace Access Capabilities (ANACAP):</b> This field indicates the capabilities associated with Asymmetric Namespace Access Reporting (refer to section 8.20).</p> <p>Bit 7 if set to '1', then the controller supports a non-zero value in the ANAGRPID field of the Namespace Management command. If cleared to '0', then the controller does not support a non-zero value in the ANAGRPID field of the Namespace Management command. If the Namespace Management command is not supported, then this bit shall be cleared to '0'.</p> <p>Bit 6 if set to '1', then the ANAGRPID field in the Identify Namespace data structure (refer to Figure 245) does not change while the namespace is attached to any controller. If cleared to '0', then the ANAGRPID field may change while the namespace is attached to any controller. Refer to section 8.20.2.</p> <p>Bit 5 is reserved.</p> <p>Bit 4 if set to '1', then the controller is able to report ANA Change state (refer to section 8.20.3.5). If cleared to '0', then the controller does not report ANA Change state.</p> <p>Bit 3 if set to '1', then the controller is able to report ANA Persistent Loss state (refer to section 8.20.3.4). If cleared to '0', then the controller does not report ANA Persistent Loss state.</p> <p>Bit 2 if set to '1', then the controller is able to report ANA Inaccessible state (refer to section 8.20.3.3). If cleared to '0', then the controller does not report ANA Inaccessible state.</p> <p>Bit 1 if set to '1', then the controller is able to report ANA Non-Optimized state (refer to section 8.20.3.2). If cleared to '0', then the controller does not report ANA Non-Optimized state.</p> <p>Bit 0 if set to '1', then the controller is able to report ANA Optimized state (refer to section 8.20.3.1). If the controller supports Asymmetric Namespace Access Reporting, then this bit is set to '1'.</p>
347:344	O	<p><b>ANA Group Identifier Maximum (ANAGRPMAX):</b> This field indicates the maximum value of a valid ANA Group Identifier for any controller in the NVM subsystem. If the controller supports Asymmetric Namespace Access Reporting (refer to the CMIC field), then this field shall be set to a non-zero value. If the controller does not support Asymmetric Namespace Access Reporting, then this field shall be cleared to 0h.</p>
351:348	O	<p><b>Number of ANA Group Identifiers (NANAGRPID):</b> This field indicates the number of ANA groups supported by this controller. If the controller supports Asymmetric Namespace Access Reporting (refer to the CMIC field), then this field shall be set to a non-zero value that is less than or equal to the ANAGRPMAX value. If the controller does not support Asymmetric Namespace Access Reporting, then this field shall be cleared to 0h.</p>
355:352	O	<p><b>Persistent Event Log Size (PELS):</b> This field indicates the maximum reportable size for the Persistent Event Log (Refer to section 5.14.1.13) in 64 KiB units. If the Persistent Event Log is not supported, then this field is reserved.</p>
511:356		Reserved

**Figure 247: Identify –Identify Controller Data Structure, I/O Command Set Independent**

Bytes	O/M <sup>1</sup>	Description
512	M	<p><b>Submission Queue Entry Size (SQES):</b> This field defines the required and maximum I/O Submission Queue entry size <del>when using the NVM Command Set.</del></p> <p>Bits 7:4 define the maximum I/O Submission Queue entry size <del>when using the NVM Command Set.</del> This value is greater than or equal to the required SQ entry size (i.e., bits 3:0 in this field). The value is in bytes and is reported as a power of two (<math>2^n</math>). The recommended value is 6, corresponding to a standard <del>NVM Command Set</del> SQ entry size of 64 bytes. Controllers that implement proprietary extensions may support a larger value.</p> <p>Bits 3:0 define the required (i.e., minimum) I/O Submission Queue <del>E</del>entry size <del>when using the NVM Command Set.</del> This is the minimum entry size that may be used. The value is in bytes and is reported as a power of two (<math>2^n</math>). The required value shall be 6, corresponding to 64.</p>
513	M	<p><b>Completion Queue Entry Size (CQES):</b> This field defines the required and maximum I/O Completion Queue entry size <del>when using the NVM Command Set.</del></p> <p>Bits 7:4 define the maximum I/O Completion Queue entry size <del>when using the NVM Command Set.</del> This value is greater than or equal to the required CQ entry size (i.e., bits 3:0 in this field). The value is in bytes and is reported as a power of two (<math>2^n</math>). The recommended value is 4, corresponding to a standard <del>NVM Command Set</del> CQ entry size of 16 bytes. Controllers that implement proprietary extensions may support a larger value.</p> <p>Bits 3:0 define the required (i.e., minimum) I/O Completion Queue entry size <del>when using the NVM Command Set.</del> This is the minimum entry size that may be used. The value is in bytes and is reported as a power of two (<math>2^n</math>). The required value shall be 4, corresponding to 16.</p>
515:514	M	<p><b>Maximum Outstanding Commands (MAXCMD):</b> Indicates the maximum number of commands that the controller processes at one time for a particular queue (which may be larger than the size of the corresponding Submission Queue). The host may use this value to size Completion Queues and optimize the number of commands submitted at one time to a particular I/O Queue. This field is mandatory for NVMe over Fabrics implementations and optional for NVMe over PCIe implementations. If the field is not used, it shall be cleared to 0h.</p>
519:516	M	<p><b>Number of Namespaces (NN):</b> This field indicates the maximum value of a valid NSID for the NVM subsystem. If the MNAN field is cleared to 0h, then this field also indicates the maximum number of namespaces supported by the NVM subsystem.</p>

**Figure 247: Identify –Identify Controller Data Structure, I/O Command Set Independent**

Bytes	O/M <sup>1</sup>	Description
521:520	M	<p><b>Optional NVM Command Support (ONCS):</b> This field indicates the optional <b>NVM I/O commands</b> and features supported by the controller. Refer to section 6.</p> <p>Bits 15:8 are reserved.</p> <p>Bit 7 if set to '1', then the controller supports the <b>NVM Command Set</b> Verify command. If cleared to '0', then the controller does not support the <b>NVM Command Set</b> Verify command.</p> <p>Bit 6 if set to '1', then the controller supports the Timestamp feature. If cleared to '0', then the controller does not support the Timestamp feature. Refer to section 5.21.1.14.</p> <p>Bit 5 if set to '1', then the controller supports reservations. If cleared to '0', then the controller does not support reservations. If the controller supports reservations then the following commands associated with reservations shall be supported: Reservation Report, Reservation Register, Reservation Acquire, and Reservation Release. Refer to section 8.8 for additional requirements.</p> <p>Bit 4 if set to '1', then the controller supports the Save field set to a non-zero value in the Set Features command and the Select field set to a non-zero value in the Get Features command. If cleared to '0', then the controller does not support the Save field set to a non-zero value in the Set Features command and the Select field set to a non-zero value in the Get Features command.</p> <p>Bit 3 if set to '1', then the controller supports the <b>NVM Command Set</b> Write Zeroes command. If cleared to '0', then the controller does not support the <b>NVM Command Set</b> Write Zeroes command.</p> <p>Bit 2 if set to '1', then the controller supports the <b>NVM Command Set</b> Dataset Management command. If cleared to '0', then the controller does not support the <b>NVM Command Set</b> Dataset Management command.</p> <p>Bit 1 if set to '1', then the controller supports the <b>NVM Command Set</b> Write Uncorrectable command. If cleared to '0', then the controller does not support the <b>NVM Command Set</b> Write Uncorrectable command.</p> <p>Bit 0 if set to '1', then the controller supports the <b>NVM Command Set</b> Compare command. If cleared to '0', then the controller does not support the <b>NVM Command Set</b> Compare command.</p>
523:522	M	<p><b>Fused Operation Support (FUSES):</b> This field indicates the fused operations that the controller supports. Refer to section 6.2.</p> <p>Bits 15:1 are reserved.</p> <p>Bit 0 if set to '1', then the controller supports the <b>NVM Command Set</b> Compare and Write fused operation. If cleared to '0', then the controller does not support the <b>NVM Command Set</b> Compare and Write fused operation. Compare shall be the first command in the sequence.</p>

**Figure 247: Identify –Identify Controller Data Structure, I/O Command Set Independent**

Bytes	O/M <sup>1</sup>	Description										
524	M	<p><b>Format NVM Attributes (FNA):</b> This field indicates attributes for the Format NVM command.</p> <p>Bits 7:34 are reserved.</p> <p>Bit 3 indicates whether the Format NVM command supports an NSID value set to FFFFFFFFh. If set to '1', then the Format NVM command does not support an NSID value set to FFFFFFFFh. If cleared to '0', then the Format NVM command supports an NSID value set to FFFFFFFFh.</p> <p>Bit 2 indicates whether cryptographic erase is supported as part of the secure erase functionality. If set to '1', then cryptographic erase is supported. If cleared to '0', then cryptographic erase is not supported.</p> <p>Bit 1 indicates whether secure erase functionality applies to all namespaces in an NVM subsystem or is specific to a particular namespace. If set to '1', then any secure erase performed as part of a format operation results in a secure erase of all namespaces in the NVM subsystem. If cleared to '0', then any secure erase performed as part of a format results in a secure erase of the particular namespace specified. If bit 3 is set to '1', then this bit shall be cleared to '0'.</p> <p>Bit 0 indicates whether the format operation (excluding secure erase) applies to all namespaces in an NVM subsystem or is specific to a particular namespace. If set to '1', then all namespaces in an NVM subsystem shall be configured with the same attributes and a format (excluding secure erase) of any namespace results in a format of all namespaces in an NVM subsystem. If cleared to '0', then the controller supports format on a per namespace basis. If bit 3 is set to '1', then this bit shall be cleared to '0'.</p>										
525	M	<p><b>Volatile Write Cache (VWC):</b> This field indicates attributes related to the presence of a volatile write cache in the controller.</p> <p>Bits 7:3 are reserved.</p> <p>Bits 2:1 indicate Flush command behavior (refer to section 6.8) if the NSID value is set to FFFFFFFFh as follows:</p> <table><tr><th>Value</th><th>Definition</th></tr><tr><td>00b</td><td>Support for the NSID field set to FFFFFFFFh is not indicated. Only controllers compliant with versions 1.3 and earlier of the specification shall be allowed to return this value.</td></tr><tr><td>01b</td><td>Reserved.</td></tr><tr><td>10b</td><td>The Flush command does not support the NSID field set to FFFFFFFFh. The controller shall fail a Flush command with the NSID set to FFFFFFFFh with a status code of Invalid Namespace or Format.</td></tr><tr><td>11b</td><td>The Flush command supports the NSID field set to FFFFFFFFh.</td></tr></table> <p>Bit 0 if set to '1' indicates that a volatile write cache is present. If cleared to '0', a volatile write cache is not present.</p> <p>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 5.21.1.6). The Flush command (refer to section 6.8) is used to request that the contents of a volatile write cache be made non-volatile.</p>	Value	Definition	00b	Support for the NSID field set to FFFFFFFFh is not indicated. Only controllers compliant with versions 1.3 and earlier of the specification shall be allowed to return this value.	01b	Reserved.	10b	The Flush command does not support the NSID field set to FFFFFFFFh. The controller shall fail a Flush command with the NSID set to FFFFFFFFh with a status code of Invalid Namespace or Format.	11b	The Flush command supports the NSID field set to FFFFFFFFh.
Value	Definition											
00b	Support for the NSID field set to FFFFFFFFh is not indicated. Only controllers compliant with versions 1.3 and earlier of the specification shall be allowed to return this value.											
01b	Reserved.											
10b	The Flush command does not support the NSID field set to FFFFFFFFh. The controller shall fail a Flush command with the NSID set to FFFFFFFFh with a status code of Invalid Namespace or Format.											
11b	The Flush command supports the NSID field set to FFFFFFFFh.											

**Figure 247: Identify –Identify Controller Data Structure, I/O Command Set Independent**

Bytes	O/M <sup>1</sup>	Description
527:526	M	<p><b>Atomic Write Unit Normal (AWUN):</b> This field indicates the size of the write operation guaranteed to be written atomically to the NVM across all namespaces with any supported namespace format during normal operation. This field is specified in logical blocks and is a 0's based value.</p> <p>If a specific namespace guarantees a larger size than is reported in this field, then this namespace specific size is reported in the NAWUN field in the Identify Namespace data structure. Refer to section 6.4.</p> <p>If a write command is submitted with size less than or equal to the AWUN value, the host is guaranteed that the write command is atomic to the NVM with respect to other read or write commands. If a write command is submitted with size greater than the AWUN value, then there is no guarantee of command atomicity. AWUN does not have any applicability to write errors caused by power failure (refer to Atomic Write Unit Power Fail).</p> <p>A value of FFFFh indicates all commands are atomic as this is the largest command size. It is recommended that implementations support a minimum of 128 KiB (appropriately scaled based on LBA size).</p> <p><b>This field is specific to the NVM Command Set and shall be cleared to 0h if the NVM Command Set is not supported.</b></p>
529:528	M	<p><b>Atomic Write Unit Power Fail (AWUPF):</b> This field indicates the size of the write operation guaranteed to be written atomically to the NVM across all namespaces with any supported namespace format during a power fail or error condition.</p> <p>If a specific namespace guarantees a larger size than is reported in this field, then this namespace specific size is reported in the NAWUPF field in the Identify Namespace data structure. Refer to section 6.4.</p> <p>This field is specified in logical blocks and is a 0's based value. The AWUPF value shall be less than or equal to the AWUN value.</p> <p>If a write command is submitted with size less than or equal to the AWUPF value, the host is guaranteed that the write is atomic to the NVM with respect to other read or write commands. If a write command is submitted that is greater than this size, there is no guarantee of command atomicity. If the write size is less than or equal to the AWUPF value and the write command fails, then subsequent read commands for the associated logical blocks shall return data from the previous successful write command. If a write command is submitted with size greater than the AWUPF value, then there is no guarantee of data returned on subsequent reads of the associated logical blocks.</p> <p><b>This field is specific to the NVM Command Set and shall be cleared to 0h if the NVM Command Set is not supported.</b></p>
530	M	<p><b>NVMI/O Command Set Vendor Specific Command Configuration (NICSVSCC):</b> This field indicates the configuration settings for <b>NVMI/O Command Set</b> Vendor Specific command handling. Refer to section 8.7.</p> <p>Bits 7:1 are reserved.</p> <p>Bit 0 if set to '1' indicates that all <b>NVMI/O Command Set</b> Vendor Specific Commands use the format defined in Figure 106. If cleared to '0' indicates that the format of all <b>NVMI/O Command Set</b> Vendor Specific Commands are vendor specific.</p>

**Figure 247: Identify –Identify Controller Data Structure, I/O Command Set Independent**

Bytes	O/M <sup>1</sup>	Description
531	M	<p><b>Namespace Write Protection Capabilities (NWPC):</b> This field indicates the optional namespace write protection capabilities supported by the controller. Refer to section 8.19.</p> <p>Bits 7:3 are reserved.</p> <p>Bit 2 if set to '1', then the controller supports the Permanent Write Protect state. If cleared to '0', then the controller does not support the Permanent Write Protect state. If this bit is set to '1', then the controller shall support the Namespace Write Protection Authentication field (refer to section 8.10).</p> <p>Bit 1 if set to '1', then the controller supports the Write Protect Until Power Cycle state. If cleared to '0', then the controller does not support Write Protect Until Power Cycle state. If this bit is set to '1', then the controller shall support the Namespace Write Protection Authentication field (refer to section 8.10).</p> <p>Bit 0 if set to '1', then the controller shall support the No Write Protect and Write Protect namespace write protection states and may support the Write Protect Until Power Cycle state and Permanent Write Protect namespace write protection states (refer to section 8.19). If cleared to '0', then the controller does not support Namespace Write Protection and bits 2:1 shall be cleared to 00b.</p>
533:532	O	<p><b>Atomic Compare &amp; Write Unit (ACWU):</b> This field indicates the size of the write operation guaranteed to be written atomically to the NVM across all namespaces with any supported namespace format for a Compare and Write fused operation.</p> <p>If a specific namespace guarantees a larger size than is reported in this field, then the Atomic Compare &amp; Write Unit size for that namespace is reported in the NACWU field in the Identify Namespace data structure. Refer to section 6.4.</p> <p>This field shall be supported if the Compare and Write fused command is supported. This field is specified in logical blocks and is a 0's based value. If a Compare and Write is submitted that requests a transfer size larger than this value, then the controller may fail the command with a status code of Invalid Field in Command. If Compare and Write is not a supported fused command, then this field shall be 0h.</p> <p><b>This field is specific to the NVM Command Set and shall be cleared to 0h if the NVM Command Set is not supported.</b></p>
535:534		Reserved



**Figure 247: Identify –Identify Controller Data Structure, I/O Command Set Independent**

Bytes	O/M <sup>1</sup>	Description																															
539:536	O	<b>SGL Support (SGLS):</b> This field indicates if SGLs are supported <del>for the NVM Command Set</del> and the particular SGL types supported. Refer to section 4.4.																															
		Bits	Description	31:22	Reserved	21	If set to '1', then the controller supports the Transport SGL Data Block descriptor. If cleared to '0', then the controller does not support the Transport SGL Data Block descriptor.	20	If set to '1', then the controller supports the Address field in SGL Data Block, SGL Segment, and SGL Last Segment descriptor types specifying an offset. If cleared to '0', then the Address field specifying an offset is not supported.	19	If set to '1', then use of a Metadata Pointer (MPTR) that contains an address of an SGL segment containing exactly one SGL Descriptor that is qword aligned is supported. If cleared to '0', then use of a MPTR containing an SGL Descriptor is not supported.	18	If set to '1', then the controller supports commands that contain a data or metadata SGL of a length larger than the amount of data to be transferred. If cleared to '0', then the SGL length shall be equal to the amount of data to be transferred.	17	If set to '1', then use of a byte aligned contiguous physical buffer of metadata (the Metadata Pointer field in Figure 105) is supported. If cleared to '0', then use of a byte aligned contiguous physical buffer of metadata is not supported.	16	If set to '1', then the SGL Bit Bucket descriptor is supported. If cleared to '0', then the SGL Bit Bucket descriptor is not supported.	15:03	Reserved	02	If set to '1', then the controller supports the Keyed SGL Data Block descriptor. If cleared to '0', then the controller does not support the Keyed SGL Data Block descriptor.	01:00	This field is used to determine the SGL support <del>for the NVM Command Set</del> . Valid values are shown in the table below.	Value	Definition	00b	SGLs are not supported.	01b	SGLs are supported. There is no alignment nor granularity requirement for Data Blocks.	10b	SGLs are supported. There is a dword alignment and granularity requirement for Data Blocks (refer to section 4.4).	11b	Reserved
		Bits	Description																														
		31:22	Reserved																														
		21	If set to '1', then the controller supports the Transport SGL Data Block descriptor. If cleared to '0', then the controller does not support the Transport SGL Data Block descriptor.																														
		20	If set to '1', then the controller supports the Address field in SGL Data Block, SGL Segment, and SGL Last Segment descriptor types specifying an offset. If cleared to '0', then the Address field specifying an offset is not supported.																														
		19	If set to '1', then use of a Metadata Pointer (MPTR) that contains an address of an SGL segment containing exactly one SGL Descriptor that is qword aligned is supported. If cleared to '0', then use of a MPTR containing an SGL Descriptor is not supported.																														
		18	If set to '1', then the controller supports commands that contain a data or metadata SGL of a length larger than the amount of data to be transferred. If cleared to '0', then the SGL length shall be equal to the amount of data to be transferred.																														
		17	If set to '1', then use of a byte aligned contiguous physical buffer of metadata (the Metadata Pointer field in Figure 105) is supported. If cleared to '0', then use of a byte aligned contiguous physical buffer of metadata is not supported.																														
		16	If set to '1', then the SGL Bit Bucket descriptor is supported. If cleared to '0', then the SGL Bit Bucket descriptor is not supported.																														
		15:03	Reserved																														
		02	If set to '1', then the controller supports the Keyed SGL Data Block descriptor. If cleared to '0', then the controller does not support the Keyed SGL Data Block descriptor.																														
		01:00	This field is used to determine the SGL support <del>for the NVM Command Set</del> . Valid values are shown in the table below.																														
Value	Definition		00b	SGLs are not supported.	01b	SGLs are supported. There is no alignment nor granularity requirement for Data Blocks.	10b	SGLs are supported. There is a dword alignment and granularity requirement for Data Blocks (refer to section 4.4).	11b	Reserved																							
Value	Definition																																
00b	SGLs are not supported.																																
01b	SGLs are supported. There is no alignment nor granularity requirement for Data Blocks.																																
10b	SGLs are supported. There is a dword alignment and granularity requirement for Data Blocks (refer to section 4.4).																																
11b	Reserved																																
543:540	O	<b>Maximum Number of Allowed Namespaces (MNAN):</b> This field indicates the maximum number of namespaces supported by the NVM subsystem. If this field is cleared to 0h, then the maximum number of namespaces supported by the NVM subsystem is less than or equal to the value in the NN field. If the controller supports Asymmetric Namespace Access Reporting, then this field shall be set to a non-zero value that is less than or equal to the NN value.																															
767:544		Reserved																															
1023:768	M	<b>NVM Subsystem NVMe Qualified Name (SUBNQN):</b> This field specifies the NVM Subsystem NVMe Qualified Name as a UTF-8 null-terminated string. Refer to section 7.9 for the definition of NVMe Qualified Name.  Support for this field is mandatory if the controller supports revision 1.2.1 or later as indicated in the Version register (refer to section 3.1.2).																															

**Figure 247: Identify –Identify Controller Data Structure, I/O Command Set Independent**

Bytes	O/M <sup>1</sup>	Description
1791:1024		Reserved
2047:1792		Refer to the NVMe over Fabrics specification.
2079:2048	M	<b>Power State 0 Descriptor (PSD0):</b> This field indicates the characteristics of power state 0. The format of this field is defined in Figure 248.
2111:2080	O	<b>Power State 1 Descriptor (PSD1):</b> This field indicates the characteristics of power state 1. The format of this field is defined in Figure 248.
2143:2112	O	<b>Power State 2 Descriptor (PSD2):</b> This field indicates the characteristics of power state 2. The format of this field is defined in Figure 248.
2175:2144	O	<b>Power State 3 Descriptor (PSD3):</b> This field indicates the characteristics of power state 3. The format of this field is defined in Figure 248.
2207:2176	O	<b>Power State 4 Descriptor (PSD4):</b> This field indicates the characteristics of power state 4. The format of this field is defined in Figure 248.
2239:2208	O	<b>Power State 5 Descriptor (PSD5):</b> This field indicates the characteristics of power state 5. The format of this field is defined in Figure 248.
2271:2240	O	<b>Power State 6 Descriptor (PSD6):</b> This field indicates the characteristics of power state 6. The format of this field is defined in Figure 248.
2303:2272	O	<b>Power State 7 Descriptor (PSD7):</b> This field indicates the characteristics of power state 7. The format of this field is defined in Figure 248.
2335:2304	O	<b>Power State 8 Descriptor (PSD8):</b> This field indicates the characteristics of power state 8. The format of this field is defined in Figure 248.
2367:2336	O	<b>Power State 9 Descriptor (PSD9):</b> This field indicates the characteristics of power state 9. The format of this field is defined in Figure 248.
2399:2368	O	<b>Power State 10 Descriptor (PSD10):</b> This field indicates the characteristics of power state 10. The format of this field is defined in Figure 248.
2431:2400	O	<b>Power State 11 Descriptor (PSD11):</b> This field indicates the characteristics of power state 11. The format of this field is defined in Figure 248.
2463:2432	O	<b>Power State 12 Descriptor (PSD12):</b> This field indicates the characteristics of power state 12. The format of this field is defined in Figure 248.
2495:2464	O	<b>Power State 13 Descriptor (PSD13):</b> This field indicates the characteristics of power state 13. The format of this field is defined in Figure 248.
2527:2496	O	<b>Power State 14 Descriptor (PSD14):</b> This field indicates the characteristics of power state 14. The format of this field is defined in Figure 248.
2559:2528	O	<b>Power State 15 Descriptor (PSD15):</b> This field indicates the characteristics of power state 15. The format of this field is defined in Figure 248.
2591:2560	O	<b>Power State 16 Descriptor (PSD16):</b> This field indicates the characteristics of power state 16. The format of this field is defined in Figure 248.
2623:2592	O	<b>Power State 17 Descriptor (PSD17):</b> This field indicates the characteristics of power state 17. The format of this field is defined in Figure 248.
2655:2624	O	<b>Power State 18 Descriptor (PSD18):</b> This field indicates the characteristics of power state 18. The format of this field is defined in Figure 248.
2687:2656	O	<b>Power State 19 Descriptor (PSD19):</b> This field indicates the characteristics of power state 19. The format of this field is defined in Figure 248.
2719:2688	O	<b>Power State 20 Descriptor (PSD20):</b> This field indicates the characteristics of power state 20. The format of this field is defined in Figure 248.
2751:2720	O	<b>Power State 21 Descriptor (PSD21):</b> This field indicates the characteristics of power state 21. The format of this field is defined in Figure 248.
2783:2752	O	<b>Power State 22 Descriptor (PSD22):</b> This field indicates the characteristics of power state 22. The format of this field is defined in Figure 248.
2815:2784	O	<b>Power State 23 Descriptor (PSD23):</b> This field indicates the characteristics of power state 23. The format of this field is defined in Figure 248.
2847:2816	O	<b>Power State 24 Descriptor (PSD24):</b> This field indicates the characteristics of power state 24. The format of this field is defined in Figure 248.
2879:2848	O	<b>Power State 25 Descriptor (PSD25):</b> This field indicates the characteristics of power state 25. The format of this field is defined in Figure 248.

**Figure 247: Identify –Identify Controller Data Structure, I/O Command Set Independent**

Bytes	O/M <sup>1</sup>	Description
2911:2880	O	<b>Power State 26 Descriptor (PSD26):</b> This field indicates the characteristics of power state 26. The format of this field is defined in Figure 248.
2943:2912	O	<b>Power State 27 Descriptor (PSD27):</b> This field indicates the characteristics of power state 27. The format of this field is defined in Figure 248.
2975:2944	O	<b>Power State 28 Descriptor (PSD28):</b> This field indicates the characteristics of power state 28. The format of this field is defined in Figure 248.
3007:2976	O	<b>Power State 29 Descriptor (PSD29):</b> This field indicates the characteristics of power state 29. The format of this field is defined in Figure 248.
3039:3008	O	<b>Power State 30 Descriptor (PSD30):</b> This field indicates the characteristics of power state 30. The format of this field is defined in Figure 248.
3071:3040	O	<b>Power State 31 Descriptor (PSD31):</b> This field indicates the characteristics of power state 31. The format of this field is defined in Figure 248.
4095:3072	O	Vendor Specific.

Figure 248 defines the power state descriptor that describes the attributes of each power state. For more information on how the power state descriptor fields are used, refer to section 8.4 on power management.

**Figure 248: Identify – Power State Descriptor Data Structure**

Bits	Description										
255:184	Reserved										
183:182	<p><b>Active Power Scale (APS):</b> This field indicates the scale for the Active Power field. If an Active Power Workload is reported for a power state, then the Active Power Scale shall also be reported for that power state.</p> <table> <tr> <th>Value</th><th>Definition</th></tr> <tr> <td>00b</td><td>Not reported for this power state</td></tr> <tr> <td>01b</td><td>0.0001 W</td></tr> <tr> <td>10b</td><td>0.01 W</td></tr> <tr> <td>11b</td><td>Reserved</td></tr> </table>	Value	Definition	00b	Not reported for this power state	01b	0.0001 W	10b	0.01 W	11b	Reserved
Value	Definition										
00b	Not reported for this power state										
01b	0.0001 W										
10b	0.01 W										
11b	Reserved										
181:179	Reserved										
178:176	<b>Active Power Workload (APW):</b> This field indicates the workload used to calculate maximum power for this power state. Refer to section 8.4.3 for more details on each of the defined workloads. This field shall not be “No Workload” unless ACTP is 0h.										
175:160	<b>Active Power (ACTP):</b> This field indicates the largest average power consumed by the NVM subsystem over a 10 second period in this power state with the workload indicated in the Active Power Workload field. The power in Watts is equal to the value in this field multiplied by the scale indicated in the Active Power Scale field. A value of 0h indicates Active Power is not reported.										
159:152	Reserved										

**Figure 248: Identify – Power State Descriptor Data Structure**

Bits	Description										
151:150	<p><b>Idle Power Scale (IPS):</b> This field indicates the scale for the Idle Power field.</p> <table> <tr> <th>Value</th><th>Definition</th></tr> <tr> <td>00b</td><td>Not reported for this power state</td></tr> <tr> <td>01b</td><td>0.0001 W</td></tr> <tr> <td>10b</td><td>0.01 W</td></tr> <tr> <td>11b</td><td>Reserved</td></tr> </table>	Value	Definition	00b	Not reported for this power state	01b	0.0001 W	10b	0.01 W	11b	Reserved
Value	Definition										
00b	Not reported for this power state										
01b	0.0001 W										
10b	0.01 W										
11b	Reserved										
149:144	Reserved										
143:128	<p><b>Idle Power (IDL P):</b> This field indicates the typical power consumed by the NVM subsystem over 30 seconds in this power state when idle (i.e., there are no pending commands, register accesses, background processes, sanitize operation, nor device self-test operations). The measurement starts after the NVM subsystem has been idle for 10 seconds. The power in Watts is equal to the value in this field multiplied by the scale indicated in the Idle Power Scale field. A value of 0h indicates Idle Power is not reported. Refer to section 8.4.</p> <p>Note: This value may be used by hosts to manage power versus resume latency. Platform and form factor specifications may have additional power measurement and reporting requirements that are outside the scope of this specification.</p>										
127:125	Reserved										
124:120	<p><b>Relative Write Latency (RWL):</b> This field indicates the relative write latency associated with this power state. The value in this field shall be less than the number of supported power states (e.g., if the controller supports 16 power states, then valid values are 0 through 15). A lower value means lower write latency.</p>										
119:117	Reserved										
116:112	<p><b>Relative Write Throughput (RWT):</b> This field indicates the relative write throughput associated with this power state. The value in this field shall be less than the number of supported power states (e.g., if the controller supports 16 power states, then valid values are 0 through 15). A lower value means higher write throughput.</p>										
111:109	Reserved										
108:104	<p><b>Relative Read Latency (RRL):</b> This field indicates the relative read latency associated with this power state. The value in this field shall be less than the number of supported power states (e.g., if the controller supports 16 power states, then valid values are 0 through 15). A lower value means lower read latency.</p>										
103:101	Reserved										
100:96	<p><b>Relative Read Throughput (RRT):</b> This field indicates the relative read throughput associated with this power state. The value in this field shall be less than the number of supported power states (e.g., if the controller supports 16 power states, then valid values are 0 through 15). A lower value means higher read throughput.</p>										
95:64	<p><b>Exit Latency (EXLAT):</b> This field indicates the maximum exit latency in microseconds associated with exiting this power state. A value of 0h indicates Exit Latency is not reported.</p>										
63:32	<p><b>Entry Latency (ENLAT):</b> This field indicates the maximum entry latency in microseconds associated with entering this power state. A value of 0h indicates Entry Latency is not reported.</p>										
31:26	Reserved										
25	<p><b>Non-Operational State (NOPS):</b> This bit indicates whether the controller processes I/O commands in this power state. If this bit is cleared to '0', then the controller processes I/O</p>										

**Figure 248: Identify – Power State Descriptor Data Structure**

Bits	Description
	commands in this power state. If this bit is set to '1', then the controller does not process I/O commands in this power state. Refer to section 8.4.1.
24	<b>Max Power Scale (MXPS):</b> This bit indicates the scale for the Maximum Power field. If this bit is cleared to '0', then the scale of the Maximum Power field is in 0.01 Watts. If this bit is set to '1', then the scale of the Maximum Power field is in 0.0001 Watts.
23:16	Reserved
15:00	<b>Maximum Power (MP):</b> This field indicates the sustained maximum power consumed by the NVM subsystem in this power state. The power in Watts is equal to the value in this field multiplied by the scale specified in the Max Power Scale bit. A value of 0h indicates Maximum Power is not reported. Refer to section 8.4.  Note: This value is intended to provide an approximate guideline for hosts to manage power versus performance. Platform and form factor specifications may have additional power measurement and reporting requirements that are outside the scope of this specification.

#### 5.15.2.3 Active Namespace ID list (CNS 02h)

A list of 1,024 namespace IDs is returned to the host containing active NSIDs in increasing order that are greater than the value specified in the Namespace Identifier (NSID) field of the command

The controller should abort the command with status code Invalid Namespace or Format if the NSID field is set to FFFFFFFEh or FFFFFFFFh. The NSID field may be cleared to 0h to retrieve a Namespace List including the namespace starting with NSID of 1h. The data structure returned is a Namespace List (refer to section 4.10)

#### 5.15.2.4 Namespace Identification Descriptor list (CNS 03h)

A list of Namespace Identification Descriptor structures (refer to Figure 249) is returned to the host for the namespace specified in the Namespace Identifier (NSID) field if it is an active NSID. If the NSID field does not specify an active NSID, then refer to section 6.1.5 for the status code to return.

The controller may return any number of variable length Namespace Identification Descriptor structures that fit into the 4,096 byte Identify payload. All remaining bytes after the namespace identification descriptor structures should be cleared to 0h, and the host shall interpret a Namespace Identifier Descriptor Length (NIDL) value of 0h as the end of the list. If, while processing these descriptors, the host detects a descriptor type that it does not recognize, then it should skip the unrecognized descriptor type and continue parsing the structure.

A controller shall not return multiple descriptors with the same Namespace Identifier Type (NIDT). A controller shall return at least one descriptor identifying the namespace (i.e., NIDT=1, NIDT=2, or NIDT=3). If bit 43 is set to '1' in the CAP register, then the Command Set Identifier NID (i.e., NIDT=4) shall be returned.

**Figure 249: Identify – Namespace Identification Descriptor**

Bytes	Description																					
00	<b>Namespace Identifier Type (NIDT):</b> This field indicates the data type contained in the Namespace Identifier field and the length of that type as defined in the following table.																					
	<table><tr><th>Value</th><th>Length (NIDL)</th><th>Definition</th></tr><tr><td>0h</td><td></td><td>Reserved</td></tr><tr><td>1h</td><td>8h</td><td><b>IEEE Extended Unique Identifier:</b> The NID field contains a copy of the EUI64 field in the Identify Namespace data structure (refer to Figure 245). If the EUI64 field of the Identify Namespace data structure is not supported, (i.e., EUI64 field is cleared to 0h), the controller shall not report a Namespace Identification Descriptor with a value of type 1h.</td></tr><tr><td>2h</td><td>10h</td><td><b>Namespace Globally Unique Identifier:</b> The NID field contains a copy of the NGUID field in the Identify Namespace data structure (refer to Figure 245). If the NGUID field of the Identify Namespace data structure is not supported (i.e., the NGUID field is cleared to 0h), the controller shall not report a Namespace Identification Descriptor with a value of type 2h.</td></tr><tr><td>3h</td><td>10h</td><td><b>Namespace UUID:</b> The NID field contains a 128-bit Universally Unique Identifier (UUID) as specified in RFC 4122. Refer to section 7.10.6.  If the namespace does not support an IEEE Extended Unique Identifier (i.e., EUI64 field is cleared to 0h) and does not support a Namespace Globally Unique Identifier (i.e., the NGUID field is cleared to 0h), then the namespace shall report a Namespace Identification Descriptor with a value of type 3h.</td></tr><tr><td>4h</td><td>1h</td><td><b>Command Set Identifier (CSI):</b> The NID field contains the I/O Command Set that operates on this namespace. Refer to Figure X1.</td></tr><tr><td>4h-5h to FFh</td><td></td><td>Reserved</td></tr></table>	Value	Length (NIDL)	Definition	0h		Reserved	1h	8h	<b>IEEE Extended Unique Identifier:</b> The NID field contains a copy of the EUI64 field in the Identify Namespace data structure (refer to Figure 245). If the EUI64 field of the Identify Namespace data structure is not supported, (i.e., EUI64 field is cleared to 0h), the controller shall not report a Namespace Identification Descriptor with a value of type 1h.	2h	10h	<b>Namespace Globally Unique Identifier:</b> The NID field contains a copy of the NGUID field in the Identify Namespace data structure (refer to Figure 245). If the NGUID field of the Identify Namespace data structure is not supported (i.e., the NGUID field is cleared to 0h), the controller shall not report a Namespace Identification Descriptor with a value of type 2h.	3h	10h	<b>Namespace UUID:</b> The NID field contains a 128-bit Universally Unique Identifier (UUID) as specified in RFC 4122. Refer to section 7.10.6.  If the namespace does not support an IEEE Extended Unique Identifier (i.e., EUI64 field is cleared to 0h) and does not support a Namespace Globally Unique Identifier (i.e., the NGUID field is cleared to 0h), then the namespace shall report a Namespace Identification Descriptor with a value of type 3h.	4h	1h	<b>Command Set Identifier (CSI):</b> The NID field contains the I/O Command Set that operates on this namespace. Refer to Figure X1.	4h-5h to FFh		Reserved
	Value	Length (NIDL)	Definition																			
	0h		Reserved																			
	1h	8h	<b>IEEE Extended Unique Identifier:</b> The NID field contains a copy of the EUI64 field in the Identify Namespace data structure (refer to Figure 245). If the EUI64 field of the Identify Namespace data structure is not supported, (i.e., EUI64 field is cleared to 0h), the controller shall not report a Namespace Identification Descriptor with a value of type 1h.																			
	2h	10h	<b>Namespace Globally Unique Identifier:</b> The NID field contains a copy of the NGUID field in the Identify Namespace data structure (refer to Figure 245). If the NGUID field of the Identify Namespace data structure is not supported (i.e., the NGUID field is cleared to 0h), the controller shall not report a Namespace Identification Descriptor with a value of type 2h.																			
	3h	10h	<b>Namespace UUID:</b> The NID field contains a 128-bit Universally Unique Identifier (UUID) as specified in RFC 4122. Refer to section 7.10.6.  If the namespace does not support an IEEE Extended Unique Identifier (i.e., EUI64 field is cleared to 0h) and does not support a Namespace Globally Unique Identifier (i.e., the NGUID field is cleared to 0h), then the namespace shall report a Namespace Identification Descriptor with a value of type 3h.																			
	4h	1h	<b>Command Set Identifier (CSI):</b> The NID field contains the I/O Command Set that operates on this namespace. Refer to Figure X1.																			
4h-5h to FFh		Reserved																				
01	<b>Namespace Identifier Length (NIDL):</b> This field contains the length in bytes of the Namespace Identifier (NID) field. The total length of the Namespace Identification Descriptor in bytes is the value in this field plus four. If this field is cleared to 0h it indicates the end of the Namespace Identifier Descriptor list.																					
02:03	Reserved																					
(NIDL + 3):04	<b>Namespace Identifier (NID):</b> This field contains a value that is globally unique and assigned to the namespace when the namespace is created. This field remains fixed throughout the life of the namespace and is preserved across namespace and controller operations (e.g., controller reset, namespace format, etc.). The type of the value is specified by the Namespace Identifier Type (NIDT) field, and the size is specified by the Namespace Identifier Length (NIDL) field.																					

#### 5.15.2.5 NVM Set List (CNS 04h)

Figure 250 defines an NVM Set List. The data structure is an ordered list by NVM Set Identifier, starting with the first NVM Set Identifier supported by the NVM subsystem that is equal to or greater than the NVM Set Identifier indicated in CDW11.NVMSETID. The NVM Set List describes the attributes for each NVM Set in the list based on the NVM Set Attributes Entry in Figure 250.

**Figure 250: NVM Set List**

Bytes	Description
00	<b>Number of Identifiers:</b> This field indicates the number of NVM Set Attributes Entries in the list. There are up to 31 entries in the list. A value of 0h indicates that there are no entries in the list.
127:01	Reserved
255:128	<b>Entry 0:</b> This field contains the first NVM Set Attributes Entry in the list, if present.
383:256	<b>Entry 1:</b> This field contains the second NVM Set Attributes Entry in the list, if present.
...	...
(N*128+255): (N*128+128)	<b>Entry N:</b> This field contains the N+1 NVM Set Attributes Entry in the list, if present.

**Figure 251: NVM Set Attributes Entry**

Bytes	Description
01:00	<b>NVM Set Identifier:</b> This field indicates the identifier of the NVM Set in the NVM subsystem that is described by this entry.
03:02	<b>Endurance Group Identifier:</b> This field indicates the Endurance Group for this NVM Set. Refer to section 8.17.
07:04	Reserved
11:08	<b>Random 4 KiB Read Typical:</b> This field indicates the typical time to complete a 4 KiB random read in 100 nanosecond units when the NVM Set is in a Predictable Latency Mode Deterministic Window and there is 1 outstanding command per NVM Set.
15:12	<b>Optimal Write Size:</b> This field indicates the size in bytes for optimal write performance. A value of 0h indicates that no Optimal Write Size is specified. This field should be set to 0h when namespaces within an NVM Set have different LBA formats that do not allow an Optimal Write Size to be specified.
31:16	<b>Total NVM Set Capacity:</b> This field indicates the total NVM capacity in this NVM Set. The value is in bytes.
47:32	<b>Unallocated NVM Set Capacity:</b> This field indicates the unallocated NVM capacity in this NVM Set. The value is in bytes.
127:48	Reserved

#### 5.15.2.6 Allocated Namespace ID list (CNS 10h)

A list of up to 1,024 namespace IDs is returned to the host containing allocated NSIDs in increasing order that are greater than the value specified in the Namespace Identifier (NSID) field of the Identify command

The controller should abort the command with status code Invalid Namespace or Format if the NSID field is set to FFFFFFFEh or FFFFFFFFh. The NSID field may be cleared to 0h to retrieve a Namespace List including the namespace starting with NSID of 1h. The data structure returned is a Namespace List (refer to section 4.10).

#### 5.15.2.7 Identify Namespace data structure for an Allocated Namespace ID (CNS 11h)

An Identify Namespace data structure (refer to [Figure 245](#) section 5.15.2.1) is returned to the host for the namespace specified in the Namespace Identifier (NSID) field if it is an allocated NSID. If the specified namespace is an unallocated NSID then the controller returns a zero filled data structure. **If the specified**



namespace is not associated with the NVM Command Set, then the controller shall abort the command with status code Invalid I/O Command Set.

If the specified namespace is an invalid NSID then the controller shall fail the command with a status code of Invalid Namespace or Format. If the NSID field is set to FFFFFFFFh then the controller should fail the command with a status code of Invalid Namespace or Format.

#### 5.15.2.8 Namespace Attached Controller list (CNS 12h)

A Controller List (refer to section 4.11) of up to 2,047 controller identifiers is returned containing a controller identifier greater than or equal to the value specified in the Controller Identifier (CDW10.CNTID) field. The list contains controller identifiers that are attached to the namespace specified in the Namespace Identifier (NSID) field. If the NSID field is set to FFFFFFFFh, then the controller should fail the command with a status code of Invalid Field in Command.

#### 5.15.2.9 Controller list (CNS 13h)

A Controller List (refer to section 4.11) of up to 2,047 controller identifiers is returned containing a controller identifier greater than or equal to the value specified in the Controller Identifier (CDW10.CNTID) field. The list contains controller identifiers in the NVM subsystem that may or may not be attached to namespace(s).

#### 5.15.2.10 Primary Controller Capabilities data structure (CNS 14h)

The Primary Controller Capabilities Structure (refer to Figure 252) is returned to the host for the primary controller specified.

**Figure 252: Identify – Primary Controller Capabilities Structure**

Bytes	Description
01:00	<b>Controller Identifier (CNTLID):</b> This field indicates the Controller Identifier of the primary controller.
...	...
79:78	<b>VI Flexible Resource Preferred Granularity (VIGRAN):</b> This field indicates the preferred granularity of assigning and removing VI Flexible Resources. Assigning and removing VI Resources in this granularity minimizes any wasted internal implementation resources.
4095:80	Reserved

#### 5.15.2.11 Secondary Controller list (CNS 15h)

A Secondary Controller List (refer to Figure 253) is returned to the host for up to 127 secondary controllers associated with the primary controller processing this command. The list contains entries for controller identifiers greater than or equal to the value specified in the Controller Identifier (CDW10.CNTID) field.

All secondary controllers are represented, including those that are in an Offline state due to SR-IOV configuration settings (e.g., VF Enable is cleared to '0' or NumVFs specifies a value that does not enable the associated secondary controller).

**Figure 253: Secondary Controller List**

Bytes	Description
00	<b>Number of Identifiers:</b> This field indicates the number of Secondary Controller Entries in the list. There are up to 127 entries in the list. A value of 0h indicates there are no entries in the list.
...	...
(N*32+63): (N*32+32)	<b>SC Entry N:</b> This field contains the N+1 Secondary Controller Entry in the list, if present.



**Figure 254: Secondary Controller Entry**

Bytes	Description
01:00	<b>Secondary Controller Identifier (SCID):</b> This field indicates the Controller Identifier of the secondary controller described by this entry.
...	...
13:12	<b>Number of VI Flexible Resources Assigned (NVI):</b> This field indicates the number of VI Flexible Resources currently assigned to the indicated secondary controller.
31:14	Reserved

#### 5.15.2.12 Namespace Granularity List (CNS 16h)

If the controller supports reporting of Namespace Granularity (refer to section 8.12.1), then a Namespace Granularity List (refer to Figure 255) is returned to the host for up to sixteen namespace granularity descriptors (refer to Figure 256).

The controller shall abort the command with status code Invalid I/O Command Set if the Command Set Identifier is not associated with the NVM Command Set.

**Figure 255: Namespace Granularity List**

Bytes	Description
03:00	<b>Namespace Granularity Attributes:</b> This field indicates attributes of the Namespace Granularity List.  Bits 31:1 are reserved.  Bit 0 ( <b>Granularity Descriptor Mapping</b> ): If set to '1', then each valid namespace granularity descriptor applies to the LBA format having the same index and the Number of Descriptors field shall be equal to the Number of LBA Formats field in the Identify Namespace data structure (refer to Figure 245). If cleared to '0', then NG Descriptor 0 shall apply to all LBA formats and the Number of Descriptors field shall be cleared to 0h.
04	<b>Number of Descriptors:</b> This field indicates the number of valid namespace granularity descriptors in the list. This is a 0's based value.  The namespace granularity descriptors with an index greater than the value in this field shall be cleared to 0h.
31:05	Reserved
47:32	<b>NG Descriptor 0:</b> This field contains the first namespace granularity descriptor in the list. This namespace granularity descriptor applies to LBA formats as indicated by the Granularity Descriptor Mapping bit.
63:48	<b>NG Descriptor 1:</b> This field contains the second namespace granularity descriptor in the list. This namespace granularity descriptor applies to LBA Format 1.
...	...
287:272	<b>NG Descriptor 15:</b> This field contains the sixteenth namespace granularity descriptor in the list. This namespace granularity descriptor applies to LBA Format 15.

The format of the namespace granularity descriptor is defined in Figure 256.

**Figure 256: Namespace Granularity Descriptor**

Bytes	Description
07:00	<b>Namespace Size Granularity:</b> Indicates the preferred granularity of allocation of namespace size when a namespace is created. The value is in bytes. A value of 0h indicates that the namespace size granularity is not reported.

**Figure 256: Namespace Granularity Descriptor**

Bytes	Description
15:08	<b>Namespace Capacity Granularity:</b> Indicates the preferred granularity of allocation of namespace capacity when a namespace is created. The value is in bytes. A value of 0h indicates that the namespace capacity granularity is not reported.

**5.15.2.13 UUID List (CNS 17h)**

The format of the UUID List is defined in Figure 257. Each UUID List entry is either 0h, the NVMe Invalid UUID, or a valid UUID. Valid UUIDs are those which are non-zero and are not the NVMe Invalid UUID (refer to section 8.24).

If bit 9 (UUID List) is set to '1' in the Controller Attributes (CTRATT) field in the Identify Controller data structure (refer to Figure 247), then:

- The UUID List shall contain at least one valid UUID (refer to section 8.24);
- The UUID 1 field shall contain a non-zero value; and
- A UUID field cleared to 0h indicates the end of the UUID List.

The list may be in any order.

**Figure 257: UUID List**

Bytes	Description
31:00	Reserved
63:32	<b>UUID 1:</b> This field contains the first UUID List Entry in the list.
95:64	<b>UUID 2:</b> This field contains the second UUID List Entry in the list, if present, otherwise cleared to 0h.
...	...
4063:4032	<b>UUID 126:</b> This field contains the last non-zero UUID List Entry in the list, if present, otherwise cleared to 0h.
4095:4064	<b>UUID 127:</b> This field shall be cleared to 0h.

The format of a UUID List Entry is defined in Figure 258.

**Figure 258: UUID List Entry**

Bytes	Description		
00	UUID Lists Entry Header:		
	Bits	Description	
	7:2	Reserved	
	1:0		
		Value	Description
		00b	No association reported.
		01b	The UUID is associated with the vendor reported in the PCI Vendor ID field of the Identify Controller data structure (refer to Figure 247).
10b		The UUID is associated with the vendor reported in the PCI Subsystem Vendor ID field of the Identify Controller data structure.	
11b	Reserved		
	Identifier Association: This field indicates whether the UUID is associated with a vendor.		
15:01	Reserved		
31:16	UUID: This field contains a 128-bit Universally Unique Identifier (UUID) as specified in RFC 4122. Refer to section 7.10.6.		

### 5.15.2.TBDa I/O Command Set specific Identify Namespace data structure (CNS 05h)

An I/O Command Set specific Identify Namespace data structure (e.g., for the NVM Command Set refer to figure X2) is returned to the host for the namespace specified in the Namespace Identifier (NSID) field if the NSID is active. If the specified namespace is an inactive NSID, then the controller returns a zero filled data structure.

The specific Identify Namespace data structure that is returned by this command is specified by the Command Set Identifier (CSI) field (refer to Figure X1). If the I/O Command Set associated with the namespace identified by the NSID field does not support the Identify Namespace data structure specified by the CSI field, the controller shall abort the command with a status code of Invalid Field in Command.

If the controller supports the Namespace Management capability (refer to section 8.12) and the NSID field is set to FFFFFFFFh, then the controller returns an I/O Command Set specific Identify Namespaces data structure that specifies capabilities that are common across namespaces for the I/O Command Set specified in the CSI field (refer to Figure X1). If the controller does not support the Namespace Management capability and the NSID field is set to FFFFFFFFh, then the controller shall fail the command with a status code of Invalid Namespace or Format.

#### 5.15.2.TBDa.1 NVM Command Set I/O Command Set Specific Identify Namespace Data structure (CSI 00h)

**Figure X2: NVM Command Set I/O Command Set Specific Identify Namespace Data Structure**

Bytes	O/M <sup>1</sup>	Description
4095:00	O	Reserved
NOTES:		
1. O/M definition: O = Optional, M = Mandatory.		

### 5.15.2.TBDb Identify I/O Command Set specific Controller data structure (CNS 06h)

An Identify Controller data structure is returned to the host for the controller processing the command. The specific Identify Controller data structure that is returned by this command is specified by the Command Set Identifier (CSI) field (refer to Figure X1). Data structures for specific I/O Command Sets are optionally defined by the I/O Command Set specifications. If the I/O Command Set specified by the CSI field does not have an Identify Controller data structure, then the controller shall return a zero filled data structure. If the host requests a data structure for an I/O Command Set that the controller does not support, the controller shall abort the command with a status code of Invalid Field in Command.

#### 5.15.2.TBDb.1 Identify I/O Command Set Specific Controller data structure (CSI 00h)

The Identify Controller data structure returned for the NVM Command Set (i.e., CSI set to 0h) is shown in figure 247.TBD.

**Figure 247.TBD: Identify –Identify Controller Data Structure, NVM Command Set**

Bytes	O/M <sup>1</sup>	Description
4095:00		Reserved

#### **5.15.2.TBDc I/O Command Set specific Active Namespace ID list (CNS 07h)**

A list of 1,024 namespace IDs is returned to the host containing active NSIDs in increasing order that are greater than the value specified in the Namespace Identifier (NSID) field of the command as specified by the Command Set Identifier (CSI) field of the command. Only namespaces associated with the I/O Command Set specified in CSI are returned. For all other CSI values the command is aborted with a status code of Invalid Field in Command.

The controller should abort the command with status code Invalid Namespace or Format if the NSID field is set to FFFFFFFEh or FFFFFFFFh. The NSID field may be cleared to 0h to retrieve a Namespace List including the namespace starting with NSID of 1h. The data structure returned is a Namespace List (refer to section 4.10)

#### **5.15.2.TBDd I/O Command Set specific Allocated Namespace ID list (CNS 1Ah)**

A list of up to 1,024 namespace IDs is returned to the host containing allocated NSIDs in increasing order that are greater than the value specified in the Namespace Identifier (NSID) field of the Identify command and as specified by the Command Set Identifier (CSI) field of the command. Only NSIDs for namespaces associated with the I/O Command Set specified in CSI are returned. For all other CSI values the command is aborted with a status code of Invalid Field in Command.

The controller should abort the command with status code Invalid Namespace or Format if the NSID field is set to FFFFFFFEh or FFFFFFFFh. The NSID field may be cleared to 0h to retrieve a Namespace List including the namespace starting with NSID of 1h. The data structure returned is a Namespace List (refer to section 4.10).

#### **5.15.2.TBDe I/O Command Set specific Identify Namespace data structure for an Allocated Namespace ID (CNS 1Bh)**

An I/O Command Set specific Identify Namespace data structure (refer to ~~Figure 245~~ section 5.15.2.TBDa) is returned to the host for the namespace specified in the Namespace Identifier (NSID) field if it is an allocated NSID. If the specified namespace is an unallocated NSID then the controller returns a zero filled data structure.

The specific Identify Namespace data structure that is returned by this command is specified by the Command Set Identifier (CSI) field in the command (refer to Figure X1). If the I/O Command Set associated with the namespace specified by the NSID field does not support the Identify Namespace data structure specified by the CSI field, the controller shall abort the command with a status code of Invalid Field in Command.

If the specified namespace is an invalid NSID then the controller shall fail the command with a status code of Invalid Namespace or Format. If the NSID field is set to FFFFFFFFh then the controller should fail the command with a status code of Invalid Namespace or Format.

#### **5.15.2.TBD Identify I/O Command Set data structure (CNS 1Ch)**

The Identify I/O Command Set data structure (refer to Figure TBDe) is returned to the host for the controller specified in the Controller ID (CNTID) field of the command if the CNTID field does not have a value of

FFFFh. If the CNTID field has a value of FFFFh, then the Identify I/O Command Set data structure is returned to the host for the controller processing the command.

This CNS value shall be implemented if bit 43 is set to '1' in the CAP register.

The Identify I/O Command Set data structure consists of an array of I/O Command Set Vectors (refer to Figure TBDb) that describe the I/O Command Sets that the controller supports and the combination of supported I/O Command Sets that may be simultaneously used. The I/O Command Set Profile Feature value indicates the index of the I/O Command Set Combination that is currently selected (refer to section 5.21.1.TBDp). I/O Command Set Combination 0 has an index value of 0h, I/O Command Set Combination 1 has an index value of 1h, and so on. Only I/O Command Sets that have a bit set to '1' in the I/O Command Set Vector of the I/O Command Set Combination selected by the I/O Command Set Profile Feature value may be used. All other I/O Command Sets are treated as unsupported I/O Command Sets.

**Figure TBDA: Identify I/O Command Set Data Structure**

Bytes	Description
7:0	<b>I/O Command Set Combination 0:</b> This field contains an I/O Command Set Vector indicating the first I/O Command Set or combination of I/O Command Sets that are simultaneously supported. If only one I/O Command Set is supported, then this field has only one bit set.
15:8	<b>I/O Command Set Combination 1:</b> This field contains an I/O Command Set Vector indicating the second I/O Command Set or combination of I/O Command Sets that are simultaneously supported if a second I/O Command Set combination is supported; otherwise, this field is cleared to 0h.  If this field is cleared to 0h, then no further I/O Command Set combinations are supported and subsequent I/O Command Set Combinations shall have a value of 0h.
23:16	<b>I/O Command Set Combination 2:</b> This field contains an I/O Command Set Vector indicating the third I/O Command Set or combination of I/O Command Sets that are simultaneously supported if a third I/O Command Set combination is supported; otherwise, this field is cleared to 0h.  If this field is cleared to 0h, then no further I/O Command Set combinations are supported and subsequent I/O Command Set Combinations shall have a value of 0h.
...	...
4095:4088	<b>I/O Command Set Combination 511:</b> This field contains an I/O Command Set Vector indicating the 511th I/O Command Set or combination of I/O Command Sets that are simultaneously supported if a 511 <sup>th</sup> I/O Command Set combination is supported; otherwise, this field is cleared to 0h.

**Figure TBDb: I/O Command Set Vector**

Bit	Description
63:1	Reserved
0	<b>NVM Command Set:</b> This bit is set to '1' if the NVM Command Set is selected. This bit is cleared to '0' if the NVM Command Set is not selected.

Editor's Note: Bit 1 – Key Value Command Set; Bit – 2 Zoned Namespace Command Set

### 5.15.3 Command Completion

Upon completion of the Identify command, the controller posts a completion queue entry to the Admin Completion Queue.

**Modify Section 5.19 Namespace Attachment command as shown below:**

### 5.19 Namespace Attachment command

The Namespace Attachment command is used to attach and detach controllers from a namespace. The attach and detach operations are persistent across all reset events.

If the Namespace Attachment command is supported, then the Namespace Management command (refer to section 20) shall also be supported.

The Namespace Attachment command uses the Data Pointer and Command Dword 10 fields. All other command specific fields are reserved.

The Select field determines the data structure used as part of the command. The data structure is 4,096 bytes in size. The data structure used for Controller Attach and Controller Detach is a Controller List (refer to section 4.11). The controllers that are to be attached or detached, respectively, are described in the data structure.

If an attempt is made to attach a namespace to a controller that does not support the corresponding I/O Command Set, then the command shall be aborted with a status code of I/O Command Set Not Supported.

If an attempt is made to attach a namespace to a controller that supports the corresponding I/O Command Set and the corresponding I/O Command Set is not enabled by the I/O Command Set profile feature, then the command shall be aborted with a status code of I/O Command Set Not Enabled.

**Figure 259: Namespace Attachment – Data Pointer**

Bits	Description
127:00	<b>Data Pointer (DPTR):</b> This field specifies the start of the data buffer. Refer to Figure 105 for the definition of this field. If using PRPs, this field shall not be a pointer to a PRP List as the data buffer may not cross more than one page boundary.

**Figure 260: Namespace Attachment – Command Dword 10**

Bits	Description								
31:04	Reserved								
03:00	<b>Select (SEL):</b> This field selects the type of attachment to perform. <table><tr><th>Value</th><th>Description</th></tr><tr><td>0h</td><td>Controller Attach</td></tr><tr><td>1h</td><td>Controller Detach</td></tr><tr><td>2h to Fh</td><td>Reserved</td></tr></table>	Value	Description	0h	Controller Attach	1h	Controller Detach	2h to Fh	Reserved
Value	Description								
0h	Controller Attach								
1h	Controller Detach								
2h to Fh	Reserved								

### 5.19.1 Command Completion

When the command is completed, the controller posts a completion queue entry to the Admin Completion Queue indicating the status for the command.

Command specific status values associated with the Namespace Attachment command are defined in Figure 261. For failures, the byte offset of the first failing entry is reported in the Command Specific Information field of the Error Information Log Entry. The controller does not process further entries in the Controller List after an error is encountered.

**Figure 261: Namespace Attachment – Command Specific Status Values**

Value	Description
18h	<b>Namespace Already Attached:</b> The controller is already attached to the namespace specified.
19h	<b>Namespace Is Private:</b> The controller is not attached to the namespace. The request to attach the controller could not be completed because the namespace is private and is already attached to one controller.
1Ah	<b>Namespace Not Attached:</b> The controller is not attached to the namespace. The request to detach the controller could not be completed.
1Ch	<b>Controller List Invalid:</b> The controller list provided is invalid or the controller list contains an administrative controller.
25h	<b>ANA Attach Failed:</b> The controller is not attached to the namespace as a result of an ANA condition (e.g., attaching the controller would result in an ANA Persistent Loss state (refer to section 8.20.3.4)).
29h	<b>I/O Command Set Not Supported:</b> The request to attach the controller could not be completed due to the I/O Command Set corresponding to the namespace is not supported by the controller.
2Ah	<b>I/O Command Set Not Enabled:</b> The request to attach the controller could not be completed due to the I/O Command Set corresponding to the namespace is restricted by the I/O Command Set profile feature.

**Modify Section 5.20 Namespace Management command as shown below:**

### 5.20 Namespace Management command

The Namespace Management command is used to manage namespaces (refer to 8.12), including create and delete operations.

Note: The controller continues to execute commands submitted to I/O Submission Queues while this operation is in progress.

If the Namespace Management command is supported, then the Namespace Attachment command (refer to section 5.19) shall also be supported.

Host software uses the Namespace Attachment command to attach or detach a namespace to or from a controller. The create operation does not attach the namespace to a controller. As a side effect of the delete operation, the namespace is detached from all controllers as the namespace is no longer present in the system. It is recommended that host software detach all controllers from a namespace prior to deleting the namespace. If the namespace is attached to another controller (i.e., a controller other than the controller processing the operation) and that controller has Namespace Attribute Notices enabled (refer to Figure 287), when a delete operation is requested, then as part of the delete operation a Namespace Attribute Notice is issued by that controller to indicate a namespace change.

The data structure used for the create operation ~~is defined in Figure 265 and has the same format as the Identify Namespace data structure defined~~ is defined by the I/O Command Set specified in the CSI field (refer to Section 5.20.x) ~~in Figure 261.. After successful completion of a Namespace Management command with the create operation, the namespace is formatted with the specified attributes. The fields that host~~

software may specify in the create operation are defined in Figure 262. Fields that are reserved are cleared to 0h by host software. There is no data structure transferred for the delete operation.

<Note to editor: remove table below>

**Figure 262: Namespace Management — Host Software Specified Fields**

Bytes	Description	Host-Specified
07:00	Namespace Size (NSZE)	Yes
15:08	Namespace Capacity (NCAP)	Yes
25:16	Reserved	
26	Formatted LBA Size (FLBAS)	Yes
28:27	Reserved	
29	End-to-end Data Protection Type Settings (DPS)	Yes
30	Namespace Multi-path I/O and Namespace Sharing Capabilities (NMIC)	Yes
91:31	Reserved	
95:92	ANA Group Identifier (ANAGRPID) <sup>4</sup>	Yes
99:96	Reserved	
101:100	NVM Set Identifier (NVMSETID)	Yes
383:102	Reserved	
Notes:		
4. A value of 0h specifies that the controller determines the value to use (refer to section 8.12.).		

The Namespace Management command uses the Data Pointer, ~~and Command Dword 10~~, and ~~Command Dword 11~~ fields. All other command specific fields are reserved.

The Namespace Identifier (NSID) field is used as follows for create and delete operations:

- **Create:** The NSID field is reserved for this operation; host software clears this field to a value of 0h. The controller shall select an available Namespace Identifier to use for the operation; or
- **Delete:** This field specifies the previously created namespace to delete in this operation. Specifying a value of FFFFFFFFh is used to delete all namespaces in the NVM subsystem. If the value of FFFFFFFFh is specified and there are zero valid namespaces, the command completes successfully.

**Figure 263: Namespace Management – Data Pointer**

Bits	Description
127:00	<b>Data Pointer (DPTR):</b> This field specifies the start of the data buffer. Refer to Figure 105 for the definition of this field. If using PRPs, this field shall not be a pointer to a PRP List as the data buffer may not cross more than one page boundary.

**Figure 264: Namespace Management – Command Dword 10**

Bits	Description
31:04	Reserved



**Figure 264: Namespace Management – Command Dword 10**

Bits	Description								
03:00	<b>Select (SEL):</b> This field selects the type of management operation to perform. <table> <tr> <th>Value</th><th>Description</th></tr> <tr> <td>0h</td><td>Create</td></tr> <tr> <td>1h</td><td>Delete</td></tr> <tr> <td>2h to Fh</td><td>Reserved</td></tr> </table>	Value	Description	0h	Create	1h	Delete	2h to Fh	Reserved
Value	Description								
0h	Create								
1h	Delete								
2h to Fh	Reserved								

**Figure 264a: Namespace Management – Command Dword 11**

Bits	Description
31:24	<b>Command Set Identifier (CSI):</b> For a create operation (i.e., SEL 0h), this field specifies the I/O Command Set for the created namespace. A CSI value of 0h (No I/O Command Set specified) creates a namespace using the NVM Command Set. For all other operations this field is reserved. Values for this field are defined by Figure X1.
23:0	Reserved

**Figure 265: Namespace Management – Data Structure for Create**

Bytes	Description
383:00	<b>Identify Namespace:</b> The fields set by host software are specified in <a href="#">section 5.20.x</a> <del>Figure 262</del> . Host software shall set reserved fields to 0h.
1023:384	Reserved
4095:1024	Vendor specific

<Note to editor: add new subsection as shown below>

## 5.20.x Host Software Specified Namespace Management Fields

The host specified namespace management fields are specific to the I/O Command Set.

### 5.20.x.1 NVM Command Set (Command Set Identifier 00h)

The data structure passed to the create operation is defined in Figure X2\_A. Fields that are reserved should be cleared to 0h by host software. This structure is a subset of the Identify Namespace structure specified for the NVM Command Set. After successful completion of a Namespace Management command with the create operation, the namespace is formatted with the specified attributes.

**Figure X2\_A: Namespace Management – Host Software Specified Fields**

Bytes	Description	Host Specified
07:00	Namespace Size (NSZE)	Yes
15:08	Namespace Capacity (NCAP)	Yes
25:16	Reserved	
26	Formatted LBA Size (FLBAS)	Yes
28:27	Reserved	

**Figure X2\_A: Namespace Management – Host Software Specified Fields**

Bytes	Description	Host Specified
29	End-to-end Data Protection Type Settings (DPS)	Yes
30	Namespace Multi-path I/O and Namespace Sharing Capabilities (NMIC)	Yes
91:31	Reserved	
95:92	ANA Group Identifier (ANAGRPID) <sup>1</sup>	Yes
99:96	Reserved	
101:100	NVM Set Identifier (NVMSETID)	Yes
383:102	Reserved	
Notes:		
1. A value of 0h specifies that the controller determines the value to use (refer to section 8.12).		

### 5.1.1 Command Completion

When the command is completed, the controller posts a completion queue entry to the Admin Completion Queue indicating the status for the command.

Namespace Management command specific status values are defined in Figure 266.

**Figure 266: Namespace Management – Command Specific Status Values**

Value	Description
0Ah	<b>Invalid Format:</b> The LBA Format specified is not supported. This may be due to various conditions, including: 1) specifying an invalid LBA Format number; 2) enabling protection information when there is not sufficient metadata per LBA; 3) the specified format is not available in the current configuration; or 4) invalid security state (refer to TCG Storage Interface Interactions specification).
15h	<b>Namespace Insufficient Capacity:</b> Creating the namespace requires more free space than is currently available. The Command Specific Information field of the Error Information Log specifies the total amount of NVM capacity required to create the namespace in bytes.
16h	<b>Namespace Identifier Unavailable:</b> The number of namespaces supported has been exceeded.
1Bh	<b>Thin Provisioning Not Supported:</b> Thin provisioning is not supported by the controller.
24h	<b>ANA Group Identifier Invalid:</b> The specified ANA Group Identifier (ANAGRPID) is not supported in the submitted command. This may be due to various conditions, including: a) specifying an ANAGRPID that does not exist; b) the controller does not allow an ANAGRPID to be specified (i.e., bit 7 in the ANACAP field is cleared to '0'); or c) the specified ANAGRPID is not supported by the controller processing the command (e.g., the specified value exceeds ANAGRPIDMAX (refer to Figure 247).  If the host specified a non-zero ANAGRPID, retrying the command with the ANAGRPID field cleared to 0h may succeed.
29h	<b>I/O Command Set Not Supported:</b> The I/O Command Set specified for a create operation is not supported by the controller.

Dword 0 of the completion queue entry contains the Namespace Identifier created. The definition of Dword 0 of the completion queue entry is in Figure 267

**Figure 267: Namespace Management – Completion Queue Entry Dword 0**

Bits	Description
31:00	<b>Namespace Identifier (NSID):</b> This field specifies the namespace identifier created in a <b>Ccreate</b> operation. This field is reserved for all other operations.

Modify Section 5.21 Set Features command as shown below:

**Figure 271: Set Features – Feature Identifiers**

Feature Identifier	Current Setting Persists Across Power Cycle and Reset <sup>2</sup>	Uses Memory Buffer for Attributes	Feature Name
00h	Reserved		
01h	No	No	Arbitration
02h	No	No	Power Management
03h	Yes	Yes	LBA Range Type
04h	No	No	Temperature Threshold
05h	No	No	Error Recovery
06h	No	No	Volatile Write Cache
07h	No	No	Number of Queues
08h	No	No	Interrupt Coalescing
09h	No	No	Interrupt Vector Configuration
0Ah	No	No	Write Atomicity Normal
0Bh	No	No	Asynchronous Event Configuration
0Ch	No	Yes	Autonomous Power State Transition
0Dh	No <sup>3</sup>	No <sup>4</sup>	Host Memory Buffer
0Eh	No	Yes	Timestamp
0Fh	No	No	Keep Alive Timer
10h	Yes	No	Host Controlled Thermal Management
11h	No	No	Non-Operational Power State Config
12h	Yes	No	Read Recovery Level Config
13h	No	Yes	Predictable Latency Mode Config
14h	No	No	Predictable Latency Mode Window
15h	No	No	LBA Status Information Report Interval
16h	No	Yes	Host Behavior Support
17h	Yes	No	Sanitize Config
18h	No	No	Endurance Group Event Configuration
19h	Yes	No	I/O Command Set Profile
1Ah to 77h	Reserved		
78h to 7Fh	Refer to the NVMe Management Interface Specification for definition.		
80h to BFh			Command Set Specific (Reserved)

**Figure 271: Set Features – Feature Identifiers**

Feature Identifier	Current Setting Persists Across Power Cycle and Reset <sup>2</sup>	Uses Memory Buffer for Attributes	Feature Name
C0h to FFh			Vendor Specific <sup>1, 5</sup>
NOTES: 1. The behavior of a controller in response to an inactive namespace ID to a vendor specific Feature Identifier is vendor specific. 2. This column is only valid if the feature is not saveable (refer to section 7.8). If the feature is saveable, then this column is not used. 3. The controller does not save settings for the Host Memory Buffer feature across power states and reset events, however, host software may restore the previous values. Refer to section 8.9. 4. The feature does not use a memory buffer for Set Features commands, but it does use a memory buffer for Get Features commands. Refer to section 8.9. 5. Selection of a UUID may be supported. Refer to section 8.24.			

**Modify Section 5.21.1 Features Specific Information as shown below:**

**5.21.1 Feature Specific Information**

Figure 271 defines the Features that may be configured with a Set Features command and retrieved with a Get Features command. ~~Figure 272 defines Features that are specific to the NVM Command Set.~~ Refer to section 7.1 for mandatory, optional, and prohibited features for the various controller types. Some Features utilize a memory buffer to configure or return attributes for a Feature, whereas others only utilize a dword in the command or completion queue entry. Feature values that are not persistent across power cycles and resets are restored to their default values as part of a controller reset operation. For more information on Features, including default value definitions, saveable value definitions, and current value definitions, refer to section 7.8.

There may be commands in execution when a Feature is changed. The new settings may or may not apply to commands already submitted for execution when the Feature is changed. Any commands submitted to a Submission Queue after a Set Features command is successfully completed shall utilize the new settings for the associated Feature. To ensure that a Features values apply to all subsequent commands, the host should allow commands being processed to complete prior to issuing the Set Features command.

If the controller does not support a changeable value for a Feature (e.g., the Feature is not changeable), and a Set Features command for that Feature is processed, then if that command specifies a Feature value that:

- is not the same as the existing value for that Feature, then the controller shall abort that command with a status code of Feature Not Changeable; and
- is the same as the existing value for that Feature, then the controller may:
  - complete that command successfully; or
  - abort that command with a status code of Feature Not Changeable.

**Figure 271: Set Features – Feature Identifiers**

Feature Identifier	Current Setting Persists Across Power Cycle and Reset <sup>2</sup>	Uses Memory Buffer for Attributes	Feature Name
00h	Reserved		
01h	No	No	Arbitration

**Figure 271: Set Features – Feature Identifiers**

Feature Identifier	Current Setting Persists Across Power Cycle and Reset <sup>2</sup>	Uses Memory Buffer for Attributes	Feature Name
02h	No	No	Power Management
03h	Yes	Yes	LBA Range Type
04h	No	No	Temperature Threshold
05h	No	No	Error Recovery
06h	No	No	Volatile Write Cache
07h	No	No	Number of Queues
08h	No	No	Interrupt Coalescing
09h	No	No	Interrupt Vector Configuration
0Ah	No	No	Write Atomicity Normal
0Bh	No	No	Asynchronous Event Configuration
0Ch	No	Yes	Autonomous Power State Transition
0Dh	No <sup>3</sup>	No <sup>4</sup>	Host Memory Buffer
0Eh	No	Yes	Timestamp
0Fh	No	No	Keep Alive Timer
10h	Yes	No	Host Controlled Thermal Management
11h	No	No	Non-Operational Power State Config
12h	Yes	No	Read Recovery Level Config
13h	No	Yes	Predictable Latency Mode Config
14h	No	No	Predictable Latency Mode Window
15h	No	No	LBA Status Information Report Interval
16h	No	Yes	Host Behavior Support
17h	Yes	No	Sanitize Config
18h	No	No	Endurance Group Event Configuration
19h to 77h	Reserved		
78h to 7Fh	Refer to the NVMe Management Interface Specification for definition.		
80h	Yes	No	Software Progress Marker
81h	No	Yes	Host Identifier
82h	No	No	Reservation Notification Mask
83h	Yes	No	Reservation Persistence
84h	No	No	Namespace Write Protection Config
85h to BFh	Reserved		
80h to BFh	Command Set Specific (Reserved)		

**Figure 271: Set Features – Feature Identifiers**

Feature Identifier	Current Setting Persists Across Power Cycle and Reset <sup>2</sup>	Uses Memory Buffer for Attributes	Feature Name
C0h to FFh			Vendor Specific <sup>1, 5</sup>
<p>NOTES:</p> <ol style="list-style-type: none"> <li>1. The behavior of a controller in response to an inactive namespace ID to a vendor specific Feature Identifier is vendor specific.</li> <li>2. This column is only valid if the feature is not saveable (refer to section 7.8). If the feature is saveable, then this column is not used.</li> <li>3. The controller does not save settings for the Host Memory Buffer feature across power states and reset events, however, host software may restore the previous values. Refer to section 8.9.</li> <li>4. The feature does not use a memory buffer for Set Features commands and does use a memory buffer for Get Features commands. Refer to section 8.9.</li> <li>5. Selection of a UUID may be supported. Refer to section 8.24.</li> </ol>			

**Figure 272: Set Features, NVM Command Set Specific—Feature Identifiers**

Feature Identifier	Current Setting Persists Across Power Cycle and Reset <sup>1</sup>	Uses Memory Buffer for Attributes	Feature Name
80h	Yes	No	Software Progress Marker
81h	No	Yes	Host Identifier
82h	No	No	Reservation Notification Mask
83h	Yes	No	Reservation Persistence
84h	No	No	Namespace Write Protection Config
85h to BFh			Reserved
<p>NOTES:</p> <ol style="list-style-type: none"> <li>1. This column is only valid if the feature is not saveable (refer to section 7.8). If the feature is saveable, then this column is not used.</li> </ol>			

*Modify header for 5.21.1.25 Software Progress Marker (Feature Identifier 80h), (Optional) – NVM Command Set Specific as shown below:*

#### 5.21.1.25 Software Progress Marker (Feature Identifier 80h), (Optional) —~~NVM Command Set Specific~~

*Modify 5.23 Format NVM command – NVM Command Set Specific as shown below:*

### 5.23 Format NVM command—~~NVM Command Set Specific~~

The Format NVM command is used to low level format the NVM media. This command is used by the host to change ~~the attributes of the NVM media (e.g., the LBA data size and/or metadata size for the NVM Command Set)~~. A low level format may destroy all data and metadata associated with all namespaces or only the specific namespace associated with the command (refer to the Format NVM Attributes field in the Identify Controller data structure). After the Format NVM command successfully completes, the controller shall not return any user data that was previously contained in an affected namespace.

As part of the Format NVM command, the host requests a format operation and may request a secure erase of the contents of the NVM (refer to the SES field in Figure 328). There are two types of secure erase. The User Data Erase erases all user content present in the NVM subsystem. The Cryptographic Erase erases all user content present in the NVM subsystem by deleting the encryption key with which the user data was previously encrypted.

The scope of the format operation and the scope of the format with secure erase depend on the attributes that the controller supports for the Format NVM command and the Namespace Identifier specified in the command as described in Figure 327. The type of secure erase, if applicable, is based on the setting of the Secure Erase Settings field in Command Dword 10 as defined in Figure 328.

**Figure 327: Format NVM – Operation Scope**

FNA Bit <sup>1</sup>	NSID	Format Operation
0	FFFFFFFFh <sup>2</sup>	All namespaces attached to the controller. Other namespaces are not affected.
0	Any allocated value (refer to 6.1.3)	Particular namespace specified. Other namespaces are not affected.
1 <sup>3</sup>	Any allocated value (refer to section 6.1.3) or FFFFFFFFFh	All namespaces 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 247) and bit 0 in the FNA field is ignored. 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 is ignored. 2. <del>If bit 3 in the FNA field is set to '1', then this value is not supported.</del> 3. <del>If bit 3 in the FNA field is set to '1', then this value does not occur. Refer to Figure 247.</del>		

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

For a Format command with the NSID field set to FFFFFFFFFh that specifies secure erase:

- a) if bit 1 is set to '1' in the FNA field (refer to Figure 247) and there are no namespaces in the NVM subsystem, then that Format command shall complete without error; and
- b) if bit 1 is cleared to '0' in the FNA field and there are no attached namespaces, then that Format command shall complete without error.

For a Format command with an NSID field set to FFFFFFFFh that does not specify a secure erase:

- a) if bit 0 is set to '1' in the FNA field and there are no namespaces in the NVM subsystem, then that Format command shall complete without error; and
- b) if bit 0 is cleared to '0' in the FNA field and there are no attached namespaces, then that Format command shall complete without error.

If bit 3 in the FNA field is set to '1' and a Format NVM command has the NSID field set to FFFFFFFFh, then the controller shall abort the command with a status code of Invalid Field In Command.

After successful completion of a Format NVM command, the settings specified in the Format NVM command (e.g., PI, MSET, LBAF) are reported as part of the Identify Namespace data structure. **For the NVM Command Set**, if the Format NVM command results in a change of the logical block size for the namespace, then the resulting namespace size (i.e., NSZE) (refer to Figure 245) and the namespace capacity (i.e., NCAP) (refer to Figure 245) may differ from the values indicated prior to the processing of the Format NVM command.

The Format NVM command uses the Command Dword 10 field. All other command specific fields are reserved.

**Figure 328: Format NVM – Command Dword 10**

Bits	Description										
31:12	Reserved										
11:09	<p><b>Secure Erase Settings (SES):</b> This field specifies whether a secure erase should be performed as part of the format and the type of the secure erase operation. The erase applies to all user data, regardless of location (e.g., within an exposed LBA, within a cache, within deallocated LBAs, etc.).</p> <table> <tr> <th>Value</th><th>Definition</th></tr> <tr> <td>000b</td><td>No secure erase operation requested</td></tr> <tr> <td>001b</td><td><b>User Data Erase:</b> All user data shall be erased, contents of the user data after the erase is indeterminate (e.g., the user data may be zero filled, one filled, etc.). The controller may perform a cryptographic erase when a User Data Erase is requested if all user data is encrypted.</td></tr> <tr> <td>010b</td><td><b>Cryptographic Erase:</b> All user data shall be erased cryptographically. This is accomplished by deleting the encryption key.</td></tr> <tr> <td>011b to 111b</td><td>Reserved</td></tr> </table>	Value	Definition	000b	No secure erase operation requested	001b	<b>User Data Erase:</b> All user data shall be erased, contents of the user data after the erase is indeterminate (e.g., the user data may be zero filled, one filled, etc.). The controller may perform a cryptographic erase when a User Data Erase is requested if all user data is encrypted.	010b	<b>Cryptographic Erase:</b> All user data shall be erased cryptographically. This is accomplished by deleting the encryption key.	011b to 111b	Reserved
Value	Definition										
000b	No secure erase operation requested										
001b	<b>User Data Erase:</b> All user data shall be erased, contents of the user data after the erase is indeterminate (e.g., the user data may be zero filled, one filled, etc.). The controller may perform a cryptographic erase when a User Data Erase is requested if all user data is encrypted.										
010b	<b>Cryptographic Erase:</b> All user data shall be erased cryptographically. This is accomplished by deleting the encryption key.										
011b to 111b	Reserved										
08	<p><b>Protection Information Location (PIL):</b> If set to '1' and protection information is enabled, then protection information is transferred as the first eight bytes of metadata. If cleared to '0' and protection information is enabled, then protection information is transferred as the last eight bytes of metadata. This setting is reported in the End-to-end Data Protection Type Settings (DPS) field of the Identify Namespace data structure and is constrained by the End-to-end Data Protection Capabilities (DPC) field of the Identify Namespace data structure.</p>										



**Figure 328: Format NVM – Command Dword 10**

Bits	Description												
07:05	<p><b>Protection Information (PI):</b> This field specifies whether end-to-end data protection is enabled and the type of protection information. The values for this field have the following meanings:</p> <table> <tr> <th>Value</th><th>Definition</th></tr> <tr> <td>000b</td><td>Protection information is not enabled</td></tr> <tr> <td>001b</td><td>Protection information is enabled, Type 1</td></tr> <tr> <td>010b</td><td>Protection information is enabled, Type 2</td></tr> <tr> <td>011b</td><td>Protection information is enabled, Type 3</td></tr> <tr> <td>100b to 111b</td><td>Reserved</td></tr> </table> <p>When end-to-end data protected is enabled, the host shall specify the appropriate protection information in the Read, Verify, Write, or Compare commands.</p>	Value	Definition	000b	Protection information is not enabled	001b	Protection information is enabled, Type 1	010b	Protection information is enabled, Type 2	011b	Protection information is enabled, Type 3	100b to 111b	Reserved
Value	Definition												
000b	Protection information is not enabled												
001b	Protection information is enabled, Type 1												
010b	Protection information is enabled, Type 2												
011b	Protection information is enabled, Type 3												
100b to 111b	Reserved												
04	<p><b>Metadata Settings (MSET):</b> This bit is set to '1' if the metadata is transferred as part of an extended data LBA. This bit is cleared to '0' if the metadata is transferred as part of a separate buffer. The metadata may include protection information, based on the Protection Information (PI) field. If the Metadata Size for the LBA Format selected is 0h, then this bit is not applicable.</p>												
03:00	<p><b>LBA Format (LBAF):</b> This field specifies the <del>LBA</del> format to apply to the NVM media. This corresponds to the <del>LBA</del> formats indicated in the Identify command, refer to <del>Figure 245 and Figure 246</del>; the Identify Namespace data structure in the applicable I/O Command Set specification. Only supported <del>LBA</del> formats shall be selected.</p>												

### 5.23.1 Command Completion

A completion queue entry is posted to the Admin Completion Queue when the NVM media format is complete. Format NVM command specific status values are defined in Figure 329.

**Figure 329: Format NVM – Command Specific Status Values**

Value	Description
0Ah	<p><b>Invalid Format:</b> The format specified is invalid. This may be due to various conditions, including:</p> <ol style="list-style-type: none"> <li>1) specifying an invalid <del>LBA</del> <del>F</del> format number;</li> <li>2) enabling protection information when there is not sufficient metadata per LBA;</li> <li>3) the specified format is not available in the current configuration; or</li> <li>4) invalid security state (refer to TCG Storage Interface Interactions specification), etc.</li> </ol>

*Modify the header for 5.24 Sanitize command as shown below:*

## 5.24 Sanitize command—~~NVM Command Set Specific~~

*Modify the header for 5.25 Security Receive command as shown below:*

## 5.25 Security Receive command—~~NVM Command Set Specific~~

*Modify the header for 5.26 Security Send command as shown below:*

## 5.26 Security Send command—~~NVM Command Set Specific~~

*Modify Section 6 NVM Command Set as shown below:*

## 6 NVM Command Set

An NVM subsystem is comprised of some number of controllers, where each controller may access some number of namespaces, where each namespace is comprised of some number of logical blocks. A logical block is the smallest unit of data that may be read or written from the controller. The logical block data size, reported in bytes, is always a power of two. Logical block sizes may be 512 bytes, 1 KiB, 2 KiB, 4 KiB, 8 KiB, etc. Supported logical block sizes are reported in the Identify Namespace data structure.

The NVM Command Set includes the commands listed in Figure 346. Refer to section 7.1 for mandatory, optional, and prohibited commands for the various controller types. The following subsections describe the definition for each of these commands. Commands shall only be submitted by the host when the controller is ready as indicated in the Controller Status register (CSTS.RDY) and after appropriate I/O Submission Queue(s) and I/O Completion Queue(s) have been created.

The Submission Queue Entry (SQE) structure and the fields that are common to all NVM commands are defined in section 4.2. The Completion Queue Entry (CQE) structure and the fields that are common to all NVM commands are defined in section 4.6. The command specific fields in the SQE and CQE structures (i.e., SQE Command Dwords 10-15, ~~and~~ CQE Dword 0, ~~and~~ CQE Dword 1) for the NVM Command Set are defined in this section.

**Figure 346: Opcodes for ~~NVM~~ I/O Commands**

Opcode by Field			Combined Opcode <sup>1</sup>	Command <sup>2</sup>	Reference <del>Section</del> <sup>5</sup>
(07)	(06:02)	(01:00)			
Standard Command	Function	Data Transfer <sup>3</sup>			
0b	000 00b	00b	00h	Flush <sup>4</sup>	<del>6.8</del>

**Figure 346: Opcodes for ~~NVM~~ I/O Commands**

0b	000 00b	01b	01h	Write	6.15
0b	000 00b	10b	02h	Read	6.9
0b	000 01b	00b	04h	Write Uncorrectable	6.16
0b	000 01b	01b	05h	Compare	6.6
0b	000 10b	00b	08h	Write Zeroes	6.17
0b	000 10b	01b	09h	Dataset Management	6.7
0b	000 11b	00b	0Ch	Verify	6.14
0b	000 11b	01b	0Dh	Reservation Register	6.11
0b	000 11b	10b	0Eh	Reservation Report	6.13
0b	001 00b	01b	11h	Reservation Acquire	6.10
0b	001 01b	01b	15h	Reservation Release	6.12
<b>Vendor Specific</b>					
1b	n/a	NOTE 3	80h to FFh	Vendor specific	
NOTES: 1. Opcodes not listed are reserved. 2. All NVM commands use the Namespace Identifier (NSID) field. The value FFFFFFFFh is not supported in this field unless footnote 4 in this figure indicates that a specific command does support that value. 3. Indicates the data transfer direction of the command. All options to the command shall transfer data as specified or transfer no data. All commands, including vendor specific commands, shall follow this convention: 00b = no data transfer; 01b = host to controller; 10b = controller to host; 11b = bidirectional. 4. This command may support the use of the Namespace Identifier (NSID) field set to FFFFFFFFh. 5. .					

**Modify Section 6.1.1 Namespace Overview as shown below:**

### 6.1.1 Namespace Overview

A namespace is a ~~formatted quantity of non-volatile memory that may be directly accessed by a host collection of logical blocks whose logical block addresses range from 0 to the size of the namespace —~~1. A namespace ID (NSID) is an identifier used by a controller to provide access to a namespace.

**Modify portions of Section 7.1.1 as shown below:**

Figure 417 to Figure 418 define commands that are mandatory, optional, and prohibited for an I/O controller.

**Figure 417: I/O Controller – Admin Command Support**

Command	Command Support Requirements <sup>1</sup>
Delete I/O Submission Queue	M

**Figure 417: I/O Controller – Admin Command Support**

Command	Command Support Requirements <sup>1</sup>
Create I/O Submission Queue	M
Get Log Page	M
Delete I/O Completion Queue	M
Create I/O Completion Queue	M
Identify	M
Abort	M
Set Features	M
Get Features	M
Asynchronous Event Request	M
Namespace Management	O
Firmware Commit	O
Firmware Image Download	O
Device Self-test	O
Namespace Attachment	O
Keep Alive	NOTE 2
Directive Send	O
Directive Receive	O
Virtualization Management	O
NVMe-MI Send	O
NVMe-MI Receive	O
Doorbell Buffer Config	O
NVMe over Fabrics Commands	Refer to the NVMe over Fabrics specification
Format NVM	O
Security Send	O
Security Send	O
Security Receive	O
Sanitize	O
I/O Command Set Specific Admin Command	Refer to the applicable I/O Command Set specification
Vendor Specific	O
Notes: 1. O = Optional, M = Mandatory, P = Prohibited 2. For NVMe over PCIe implementations, the Keep Alive command is optional. For NVMe over Fabrics implementations, the associated NVMe Transport binding defines whether the Keep Alive command is optional or mandatory.	

**Figure 418: I/O Controller – NVM Command Set Specific Admin Command Support**

Command	Command Support Requirements <sup>1</sup>
Format NVM	⊖
Security Send	⊖
Security Send	⊖
Security Receive	⊖
Sanitize	⊖
Get LBA Status	O
Notes: 1. O = Optional, M = Mandatory, P = Prohibited	

**Figure 419: I/O Controller – NVM Command Set Support**

Command	Command Support Requirements <sup>1</sup>
Flush	M
Write	M
Read	M
Write Uncorrectable	O
Compare	O
Write Zeroes	O
Dataset Management	O
Verify	O
Reservation Register	O <sup>2</sup>
Reservation Report	O <sup>2</sup>
Reservation Acquire	O <sup>2</sup>
Reservation Release	O <sup>2</sup>
Vendor Specific	O
Notes: 1. O = Optional, M = Mandatory, P = Prohibited 2. Mandatory if reservations are supported as indicated in the Identify Controller data structure.	

Figure 420 and Figure 421 define log pages that are mandatory, optional, and prohibited for an I/O controller.

**Figure 420: I/O Controller – Log Page Support**

Log Page Name	Command Support Requirements <sup>1</sup>
Error Information	M

**Figure 420: I/O Controller – Log Page Support**

Log Page Name	Command Support Requirements <sup>1</sup>
SMART / Health Information (Controller scope)	M
SMART / Health Information (NVM subsystem scope)	O
Firmware Slot Information	M
Changed Namespace List	O
Commands Supported and Effects	O
Device Self-test	O
Telemetry Host-Initiated	O
Telemetry Controller-Initiated	O
Endurance Group Information	O
Predictable Latency Per NVM Set	O
Predictable Latency Event Aggregate	O
Asymmetric Namespace Access	O
Persistent Event	O
LBA Status Information	O
Endurance Group Event Aggregate	O
Notes: 1. O = Optional, M = Mandatory, P = Prohibited	

**Figure 421: I/O Controller – NVM Command Set Specific Log Page Support**

Log Page Name	Command Support Requirements <sup>1</sup>
Reservation Notification	O
Sanitize Status	O
Notes: 1. O = Optional, M = Mandatory, P = Prohibited	

Figure 422 and Figure 423 define features that are mandatory, optional, and prohibited for an I/O controller.

**Figure 422: I/O Controller – Feature Support**

Feature Name	Feature Support Requirements <sup>1</sup>	Logged in Persistent Event Log <sup>1</sup>
Arbitration	M	O
Power Management	M	NR
LBA Range Type	O	NR
Temperature Threshold	M	O

**Figure 422: I/O Controller – Feature Support**

Feature Name	Feature Support Requirements <sup>1</sup>	Logged in Persistent Event Log <sup>1</sup>
Error Recovery	M	O
Volatile Write Cache	O	O
Number of Queues	M	O
Interrupt Coalescing	NOTE 2	O
Interrupt Vector Configuration	NOTE 2	O
Write Atomicity Normal	M	O
Asynchronous Event Configuration	M	NR
Autonomous Power State Transition	O	O
Host Memory Buffer	O	O
Timestamp	O	P
Keep Alive Timer	O	O
Host Controlled Thermal Management	O	O
Non-Operational Power State Config	O	O
Read Recovery Level Config	O	O
Predictable Latency Mode Config	O	O
Predictable Latency Mode Window	O	P
LBA Status Information Attributes	O	O
Host Behavior Support	O	O
Sanitize Config	O	O
Endurance Group Event Configuration	O	O
<b>I/O Command Set Profile</b>	<b>O</b>	<b>O</b>
Notes: 1. O = Optional, M = Mandatory, P = Prohibited, NR = Not Recommended 2. The feature is mandatory for NVMe over PCIe. This feature is not supported for NVMe over Fabrics.		

**Figure 423: I/O Controller – NVM Command Set Specific Feature Support**

Feature Name	Feature Support Requirements <sup>1</sup>	Logged in Persistent Event Log <sup>1</sup>
Software Progress Marker	O	NR
Host Identifier	O <sup>2</sup>	O
Reservation Notification Mask	O <sup>3</sup>	O
Reservation Persistence	O <sup>3</sup>	O

**Figure 423: I/O Controller – NVM Command Set Specific Feature Support**

Feature Name	Feature Support Requirements <sup>1</sup>	Logged in Persistent Event Log <sup>1</sup>
Namespace Write Protection Config	O	O
Notes: 1. O = Optional, M = Mandatory, P = Prohibited, NR = Not Recommended 2. Mandatory if reservations are supported as indicated in the Identify Controller data structure. 3. Mandatory if reservations are supported by the namespace as indicated by a non-zero value in the Reservation Capabilities (RESCAP) field in the Identify Namespace data structure.		

*Modify portions of Section 7.1.2 as shown below:*

### 7.1.2 Administrative Controller

...

Figure 426 and Figure 427 define commands that are mandatory, optional, and prohibited for an administrative controller. Since an administrative controller does not support I/O queues, NVM Command Set commands that are not admin commands are not supported. A host may utilize the Commands Supported and Effects log page to determine optional commands that are supported by an Administrative controller.

**Figure 426: Administrative Controller – Admin Command Support**

Command	Command Support Requirements <sup>1</sup>
Delete I/O Submission Queue	P
Create I/O Submission Queue	P
Get Log Page	M
Delete I/O Completion Queue	P
Create I/O Completion Queue	P
Identify	M
Abort	O
Set Features	O <sup>3</sup>
Get Features	O <sup>3</sup>
Asynchronous Event Request	O <sup>4</sup>
Namespace Management	O
Firmware Commit	O
Firmware Image Download	O
Device Self-test	O
Namespace Attachment	O
Keep Alive	NOTE 2
Directive Send	O
Directive Receive	O
Virtualization Management	O
NVMe-MI Send	O
NVMe-MI Receive	O
Doorbell Buffer Config	O
NVMe over Fabrics Commands	Refer to the NVMe over Fabrics specification
Format NVM	O
Security Send	O
Security Receive	O



**Figure 426: Administrative Controller – Admin Command Support**

Command	Command Support Requirements <sup>1</sup>
Sanitize	O
I/O Command Set Specific Admin Command	⊖ Refer to the applicable I/O Command Set specification
Vendor Specific	O
Notes: 1. O = Optional, M = Mandatory, P = Prohibited 2. For NVMe over PCIe implementations, the Keep Alive command is optional. For NVMe over Fabrics implementations, the associated NVMe Transport binding defines whether the Keep Alive command is optional or mandatory. 3. Mandatory if any of the features in Figure 430 are implemented. 4. Mandatory if Telemetry Log, Firmware Commit, or SMART/Health Critical Warnings are supported.	

**Figure 427: Administrative Controller – NVM Command Set Specific Admin Command Support**

Command	Command Support Requirements <sup>1</sup>
Format NVM	⊖
Security Send	⊖
Security Receive	⊖
Sanitize	⊖
Get LBA Status	P
Notes: 1. O = Optional, M = Mandatory, P = Prohibited	

...

Figure 430 and Figure 431 defines features that are mandatory, optional, and prohibited for an administrative controller. If any feature is supported, then the Set Features and Get Features commands shall be supported.

**Figure 430: Administrative Controller – Feature Support**

Feature Name	Feature Support Requirements <sup>1</sup>	Logged in Persistent Event Log <sup>1</sup>
Arbitration	P	P
Power Management	O	NR
LBA Range Type	P	P
Temperature Threshold	O	O
Error Recovery	P	P
Volatile Write Cache	P	P
Number of Queues	P	P
Interrupt Coalescing	NOTE 2	NOTE 2
Interrupt Vector Configuration	NOTE 2	NOTE 2
Write Atomicity Normal	P	P
Asynchronous Event Configuration	O <sup>3</sup>	NR
Autonomous Power State Transition	O	O
Host Memory Buffer	O	O
Timestamp	O	P
Keep Alive Timer	O	O
Host Controlled Thermal Management	O	O
Non-Operational Power State Config	O	O

**Figure 430: Administrative Controller – Feature Support**

Feature Name	Feature Support Requirements <sup>1</sup>	Logged in Persistent Event Log <sup>1</sup>
Read Recovery Level Config	O	O
Predictable Latency Mode Config	O	P
Predictable Latency Mode Window	O	O
LBA Status Information Attributes	P	O
Host Behavior Support	O	O
Sanitize Config	O	O
Endurance Group Event Configuration	O	O
<b>I/O Command Set Profile</b>	<b>P</b>	<b>P</b>
Notes: 1. O = Optional, M = Mandatory, P = Prohibited, NR = Not Recommended 2. The feature is optional for NVMe over PCIe. This feature is not supported for NVMe over Fabrics. Mandatory if Telemetry Log, Firmware Commit or SMART/Health Critical Warnings are supported.		

*Modify 8.8.3 figure 462 Command Behavior in the Presence of a Reservation as shown below:*

**Figure 462: Command Behavior in the Presence of a Reservation**

NVMe Command	Write Exclusive Reservation		Exclusive Access Reservation		Write Exclusive Registrants Only or Write Exclusive All Registrants Reservation		Exclusive Access Registrants Only or Exclusive Access All Registrants Reservation	
	Non-Registrant	Registrant	Non-Registrant	Registrant	Non-Registrant	Registrant	Non-Registrant	Registrant
<b><del>NVM</del> Read Command Group</b>								
Compare <sup>3</sup>								
Read <sup>3</sup>								
Security Receive (Admin)	A	A	C	C	A	A	C	A
Verify <sup>3</sup>								
<b>I/O Command Set Specific Read commands<sup>2</sup></b>								
<b><del>NVM</del> Write Command Group</b>								

Figure 462: Command Behavior in the Presence of a Reservation

NVMe Command	Write Exclusive Reservation		Exclusive Access Reservation		Write Exclusive Registrants Only or Write Exclusive All Registrants Reservation		Exclusive Access Registrants Only or Exclusive Access All Registrants Reservation	
	Non-Registrant	Registrant	Non-Registrant	Registrant	Non-Registrant	Registrant	Non-Registrant	Registrant
Dataset Management <sup>3-</sup> Flush Format NVM (Admin) Namespace Attachment (Admin) Namespace Management (Admin) Sanitize (Admin) Security Send (Admin) Write <sup>3</sup> Write Uncorrectable <sup>3</sup> Write Zeroes <sup>3</sup> <b>I/O Command Set Specific Write commands<sup>2</sup></b>	C	C	C	C	C	A	C	A
<b>Reservation Command Groups</b>								
Reservation Acquire - Acquire	C	C	C	C	C	C	C	C
Reservation Acquire - Preempt Reservation Acquire - Preempt and Abort Reservation Release	C	A	C	A	C	A	C	A
<b>All Other Commands Group</b>								
All other commands <sup>1</sup>	A	A	A	A	A	A	A	A

**Figure 462: Command Behavior in the Presence of a Reservation**

NVMe Command	Write Exclusive Reservation		Exclusive Access Reservation		Write Exclusive Registrants Only or Write Exclusive All Registrants Reservation		Exclusive Access Registrants Only or Exclusive Access All Registrants Reservation	
	Non-Registrant	Registrant	Non-Registrant	Registrant	Non-Registrant	Registrant	Non-Registrant	Registrant
<p>Key:</p> <p>A definition: A=Allowed, command processed normally by the controller</p> <p>C definition: C=Conflict, command aborted by the controller with status Reservation Conflict</p> <p>Notes:</p> <p>1. The behavior of a vendor specific command is vendor specific.</p> <p>2. For I/O Command Sets other than the NVM Command Set, refer to the I/O Command Set specification</p> <p>3. For NVM Command Set</p>								

**Modify Section 8.12 Namespace Management (optional) as shown below:**

The size of a namespace is based on the ~~number of logical blocks size~~ requested in a create operation, the format of the namespace, and any characteristics (e.g., endurance). The controller determines the NVM capacity allocated for that namespace. Namespaces may be created with different usage characteristics (e.g., endurance) that utilize differing amounts of NVM capacity. Namespace characteristics and the mapping of these characteristics to NVM capacity usage are outside the scope of this specification.

The total and unallocated NVM capacity for the NVM subsystem is reported in the Identify Controller data structure (refer to Figure 247). For controllers that support NVM Sets, the total and unallocated NVM capacity for each NVM Set is reported as part of the NVM Set Attributes Entry (refer to Figure 251). For each namespace, the NVM Set in which the namespace is allocated is 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 the namespace.

For each namespace, the NVM capacity used for that namespace is reported in the Identify Namespace data structure (refer to Figure 245). The controller may allocate NVM capacity in units such that the requested size for a namespace may be rounded up to the next unit boundary. The units in which NVM capacity is allocated are reported in the Namespace Granularity List (refer to Figure 255) if supported. For example, if host software requests a namespace of 32 logical blocks with a logical block size of 4 KiB for a total size of 128 KiB and the allocation unit for the implementation is 1 MiB, then the NVM capacity consumed may be rounded up to 1 MiB. The NVM capacity fields may not correspond to the logical block size multiplied by the total number of logical blocks.

## NVM Express Management Interface Specification 1.1 Changes

Modify Section 5.20 Namespace Management command as shown below:

### 5.7 Read NVMe-MI Data Structure

The Read NVMe-MI Data Structure command requests data that describes information about the NVM Subsystem, the Management Endpoint, or the NVMe Controllers.

The command uses NVMe Management Dword 0 and **Dword 1**. The format of NVMe Management Dword 0 is shown in Figure 88 and the Format of NVMe Management Dword 1 is shown in Figure 88a. ~~NVMe Management Dword 1 is reserved.~~ There is no Request Data included in a Read NVMe-MI Data Structure command.

Figure 88: Read NVMe-MI Data Structure – NVMe Management Dword 0

Bit	Description																
31:24	<b>Data Structure Type (DTYP):</b> This field specifies the data structure to return. <table><tr><th>Value</th><th>Definition</th></tr><tr><td>00h</td><td>NVM Subsystem Information</td></tr><tr><td>01h</td><td>Port Information</td></tr><tr><td>02h</td><td>Controller List</td></tr><tr><td>03h</td><td>Controller Information</td></tr><tr><td>04h</td><td>Optional Commands Supported</td></tr><tr><td>05h</td><td>Management Endpoint Buffer Command Support List</td></tr><tr><td>06h to FFh</td><td>Reserved</td></tr></table>	Value	Definition	00h	NVM Subsystem Information	01h	Port Information	02h	Controller List	03h	Controller Information	04h	Optional Commands Supported	05h	Management Endpoint Buffer Command Support List	06h to FFh	Reserved
Value	Definition																
00h	NVM Subsystem Information																
01h	Port Information																
02h	Controller List																
03h	Controller Information																
04h	Optional Commands Supported																
05h	Management Endpoint Buffer Command Support List																
06h to FFh	Reserved																
23:16	<b>Port Identifier (PORTID):</b> This field contains the identifier of the port whose data structure is returned. If the DTYP field value corresponds to Port Information, then this field contains the Port Identifier whose information is requested. If the DTYP field value corresponds to Management Endpoint Buffer Command Support List, then this field contains the Port Identifier whose information is requested. For all other values of the DTYP field, this field is reserved.																
15:00	<b>Controller Identifier (CTRLID):</b> This field contains the Controller identifier whose data structure is returned. If the DTYP field value corresponds to Controller List (i.e., 02h), <del>or</del> Controller Information (i.e., 03h), or Optional Commands Supported (i.e., 04h), then this field contains the Controller identifier in the NVM Subsystem whose information is requested. For all other values of the DTYP field, this field is reserved.																

Figure 88a: Read NVMe-MI Data Structure – NVMe Management Dword 1

Bit	Description
31:08	Reserved
07:00	<b>I/O Command Set Identifier (IOCSI):</b> If the DTYP field value corresponds to Optional Commands Supported or the Management Endpoint Buffer Command Support List, then this field specifies the I/O Command Set

**Figure 88a: Read NVMe-MI Data Structure – NVMe Management Dword 1**

Bit	Description
	used to select the optional I/O Command Set Specific Admin commands. For more information about I/O Command Sets refer to the NVMe Express specification.
	For all other values of the DTYP field, this field is reserved.

...

The Optionally Supported Command List data structure contains a list of optional commands that a Responder supports. The I/O Command Set Identifier (IOCSI) field in NVMe Management Dword 1 selects the I/O Command Set for the I/O Command Set Specific Admin commands that are returned in the Optionally Supported Command List data structure. The Optionally Supported Command List data structure may contain up to 2,047 commands, and shall be minimally sized (i.e., if there is one optionally supported command, the data structure is 4 bytes total).

**Figure 95: Optionally Supported Command List Data Structure**

Byte	Description
01:00	<b>Number of Commands (NUMCMD):</b> This field contains the number of optionally supported commands in the list. A value of 0h indicates there are no commands in the list.
03:02	<b>Command 0 (CMD0):</b> This field contains the Command Type and Opcode for the first optionally supported command or 0h if the list is empty (i.e., no optional commands are supported). Refer to Figure 96.
05:04	<b>Command 1 (CMD1):</b> This field contains the Command Type and Opcode for the second optionally supported command, if applicable. Refer to Figure 96.
...	
(N*2 +3): (N*2 + 2)	<b>Command N (CMDN):</b> This field contains the Command Type and Opcode for the N+1 optionally supported command, if applicable. Refer to Figure 96.

**Figure 96: Optionally Supported Command Data Structure**

Byte	Description		
00	<b>Command Type:</b> This field specifies the command set used by the optionally supported command.		
		<b>Bits</b>	<b>Description</b>
		7	Reserved
		6:3	<b>NVMe-MI Message Type (NMIMT):</b> This field specifies the NVMe-MI Message Type. Refer to Figure 18.
		2:0	Reserved
01	<b>Opcode:</b> This field specifies the opcode used for the optionally supported command.		

If the Management Endpoint Buffer Size field in the Port Information Data Structure is not 0h, then returning of the Management Endpoint Buffer Command Support List data structure shall be supported by the Management Endpoint. If the Management Endpoint Buffer Size field in the Port Information Data Structure is 0h, then the Data Structure Type value for Management Endpoint Buffer Command Support List is reserved.

The Management Endpoint Buffer Command Support List data structure contains a list of commands that support the use of the Management Endpoint Buffer. The I/O Command Set Identifier (IOCSI) field in NVMe Management Dword 1 selects the I/O Command Set for the I/O Command Set Specific Admin commands

that are returned in the Management Endpoint Buffer Command Support List data structure. The data structure may contain up to 2,047 commands, and shall be minimally sized (i.e., if there is 1 optionally supported command, the data structure is 4 bytes total).

The list of commands that support the Management Endpoint Buffer may be different among Management Endpoints within the NVM Subsystem. The Port Identifier (PORTID) field in NVMe Management Dword 0 of the Read NVMe-MI Data Structure specifies the port of the Management Endpoint whose Management Endpoint Buffer Command Support List data structure is returned.

**Figure 97: Management Endpoint Buffer Supported Command List Data Structure**

Byte	Description
01:00	<b>Number of Commands (NUMCMD):</b> This field contains the number of commands in the list. A value of 0h indicates there are no commands in the list.
03:02	<b>Command 0 (CMD0):</b> This field contains the Management Endpoint Buffer Supported Command Data Structure (refer to Figure 98) for the first command that supports the use of the Management Endpoint Buffer associated with the Management Endpoint.
05:04	<b>Command 1 (CMD1):</b> This field contains the Management Endpoint Buffer Supported Command Data Structure (refer to Figure 98) for the second command that supports the use of the Management Endpoint Buffer associated with the Management Endpoint.
...	
(N*2 + 3): (N*2 + 2)	<b>Command N (CMDN):</b> This field contains the Management Endpoint Buffer Supported Command Data Structure (refer to Figure 98) for the N+1 command that supports the use of the Management Endpoint Buffer associated with the Management Endpoint.

**Figure 98: Management Endpoint Buffer Supported Command Data Structure**

Byte	Description								
00	<b>Command Type:</b> This field specifies the <b>command set</b> that supports the Management Endpoint Buffer. <table border="1"> <tr> <th>Bits</th><th>Description</th></tr> <tr> <td>7</td><td>Reserved</td></tr> <tr> <td>6:3</td><td><b>NVMe-MI Message Type (NMIMT):</b> This field specifies the NVMe-MI Message Type. Refer to Figure 18.</td></tr> <tr> <td>2:0</td><td>Reserved</td></tr> </table>	Bits	Description	7	Reserved	6:3	<b>NVMe-MI Message Type (NMIMT):</b> This field specifies the NVMe-MI Message Type. Refer to Figure 18.	2:0	Reserved
Bits	Description								
7	Reserved								
6:3	<b>NVMe-MI Message Type (NMIMT):</b> This field specifies the NVMe-MI Message Type. Refer to Figure 18.								
2:0	Reserved								
01	<b>Opcode:</b> This field specifies the opcode of the command that supports the Management Endpoint Buffer.								