



**LEGAL NOTICE:**

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

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

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

**NOTICE TO NON-MEMBERS OF NVM EXPRESS, INC.:** If you are not a Member of NVM Express, Inc. and you have obtained a copy of this document, you only have a right to review this document or make reference to or cite this document. Any such references or citations to this document must acknowledge NVM Express, Inc. copyright ownership of this document. The proper copyright citation or reference is as follows: "© 2008 to 2024 NVM Express, Inc. ALL RIGHTS RESERVED." When making any such citations or references to this document you are not permitted to revise, alter, modify, make any derivatives of, or otherwise amend the referenced portion of this document in any way without the prior express written permission of NVM Express, Inc. Nothing contained in this document shall be deemed as granting you any kind of license to implement or use this document or the specification described therein, or any of its contents, either expressly or impliedly, or to any intellectual property owned or controlled by NVM Express, Inc., including, without limitation, any trademarks of NVM Express, Inc.

**LEGAL DISCLAIMER:**

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

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

The NVM Express® design mark is a registered trademark of NVM Express, Inc. PCI-SIG®, PCI Express®, and PCIe® are registered trademarks of PCI-SIG. InfiniBand™ is a trademark and servicemark of the InfiniBand Trade Association.

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

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

## NVM Express® Technical Errata

<b>Errata ID</b>	<b>118</b>
<b>Revision Date</b>	<b>1/8/2024</b>
<b>Affected Spec Ver.</b>	<b>NVM Express® Base Specification Revision 2.0c</b> <b>NVM Express® NVM Command Set Specification Revision 1.0c</b> <b>NVM Express® Management Interface Specification Revision 1.2c</b> <b>NVM Express® TCP Transport Specification 1.0c</b> <b>NVM Express® Key Value Command Set Specification 1.0c</b>
<b>References</b>	<b>ECN115</b>

### Errata Author(s)

<b>Name</b>	<b>Company</b>
Judy Brock, Mike Allison, Bill Martin,	Samsung
Fred Knight	NetApp
David Black, Austin Bolen	DellEMC
Paul Suhler	Kioxia
Andres Baez, Myron Loewen	Solidigm
Randy Jennings	Pure Storage
Matthew Goepfert	HPE

### Errata Overview

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

## Revision History

Revision Date	Change Description
5/24/2023	Initial creation with the following bug solutions: <ul style="list-style-type: none"> <li>Bug 130 Additional clarification text to the Persistent Event Header Return Size that can Cause Data Corruption.</li> <li>Bug 11 Remove controllers attached to namespaces</li> <li>Bug 21 invalid EILBRT and ILBRT values</li> </ul>
5/31/2023	Incorporated the following bug solutions: <ul style="list-style-type: none"> <li>Bug 98 Fabrics Command Capsule SQE Corrections</li> <li>Bug 123 Capacity Exceeded Description and Use Are Broken</li> <li>Bug 135 Incorrect section reference in Figure 94</li> <li>Bug 9 Correct namespace for NSID usage</li> </ul>
6/7/2023	Incorporated the following bug solutions: <ul style="list-style-type: none"> <li>Bug 83 Unsafe shutdown clarification</li> <li>Bug 68 Type, byte instead of bit</li> <li>Bug 10 Consistency of "namespace identifier", "Namespace ID"</li> </ul>
6/14/2023	Incorporated the following bug solutions: <ul style="list-style-type: none"> <li>Bug 142 The hyphen in NVMe-MI should be non-breaking</li> <li>Bug 141 Bad Hyperlink to Reservation Acquire command</li> </ul>
6/21/2023	Incorporated the following bug solutions: <ul style="list-style-type: none"> <li>Bug 7 Figure 7: TCP handshake is incomplete</li> </ul>
6/28/2023	Incorporated the following bug solutions: <ul style="list-style-type: none"> <li>Bug 70 PRINFO applies to both Read and Write portion of Compare Command</li> <li>Bug 35 Fabrics: Dynamic controllers - initially attached namespaces</li> <li>Bug 91 Need a blunter UTF-8 warning that NQNs are not IQNs</li> <li>Bug 101 Clarify "Not Required" value for secure channel in Discovery Log Page</li> </ul>
7/12/2023	Incorporated the following bug solutions: <ul style="list-style-type: none"> <li>Bug 12 Incorrect statement of changing FID</li> <li>Bug 47 Delete reference to SMBus and VDM providing same functionality</li> <li>Bug 76 CIAP and MEB Error Checking Issue</li> <li>Bug 110 Revise KV NUSE definition</li> <li>Bug 66 Keep Alive Support for Discovery Controllers</li> </ul>
7/14/2023	Cleaned up ECN to get ready for member review.
7/17/2023	Removed the redundant Figure 19 in the NVMe-MI spec. Changed the row properties in remaining Figure 19 of the NVMe-MI spec to get the figure header to properly continue on each page.
7/19/2023	Updated Capacity Exceeded to remove specific I/O Command Sets (as a result of solution for Bug 138). Fixed the navigation pane to all nesting is shown and removed repeated section with no changes.
7/20/2023	Editorial changes during NVM Express Technical Workgroup review.
7/27/2023	Editorial changes during NVM Express Technical Workgroup review. Highlighted modified or inserted text that is a figure reference or a section reference that was not already highlighted.
9/6/2023	Resolving comments submitted by Randy Jennings.
9/9/2023	Resolving comments submitted by Randy Jennings.
9/27/2023	Errata Task Group edits when resolving Randy Jennings open issues. Deleted section 2.3.3 in the Base Specification since no text changes were identified, the review comments were moved to bug 209.
10/10/2023	Editorial changes from Randy Jennings.
10/26/2023	Updates to description per TWG review. Accepted all changes.
11/17/2023	Integrated
12/22/2023	Updated per Mike Allison and Fred Knight Review

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

Revision Date	Change Description
1/8/2024	Updated to 2024. Converted all references to text.

# Description of Changes

## NVM Express Base Specification 2.0c:

### Backward Incompatible Changes:

- The definition of the Unsafe Shutdown field in the SMART / Health Information log page is updated to include the host shutting down controllers using the CC.SHN property as the previous text only included the NVM Subsystem Shutdown.

### Editorial Changes:

- Clarified that the host is responsible for allocating a buffer of at least 512 bytes for the Persistent Event log page action of Establish Context and Read 512 Bytes of Header.
- Clarify that namespaces are attached to controllers and not controllers attached to namespaces.
- Moved Bytes 03:00 from Fabrics Command Capsule SQE Format figure into new Fabrics Command Capsule Command Dword 0 figure and state Fabrics-specific requirements
- Added SGL1 field to Fabrics Command Capsule SQE Format figure and have Authentication Receive, Authentication Send, and Connect commands reference it.
- Collapsed Bytes 03:00 in each Fabrics command figure and refer to the new Fabrics Command Capsule Command Dword 0 figure.
- Clarified in Property Get, Property Set, and Disconnect commands that Byte 01 is cleared to 0h.
- Clarified the definition of the Capacity Exceeded status code..
- Corrected the referenced section in the definition of the Invalid Use of Controller Memory Buffer status code.
- Clarified the comparison of the specified NSID as opposed to the specified namespace when comparing to an inactive NSID, an allocated NSID, an unallocated NSID, or an invalid NSID.
- Corrected the capitalization of namespace ID for consistency with the definition.
- Changed from the use of “namespace ID” to the use of “NSID” in sections that mixed the usage.
- Corrected reference to Reservation Acquire command
- Use a non-breaking hyphen in “NVMe-MI”.
- Explicitly state that initially attached namespaces for a dynamic controller may vary by controller among controllers for the same host. This is already the case via the ability of the NVM subsystem to autonomously change namespace attachment after controller creation.
- Expanded details of UTF-8 support by adding improved description of UTF-8 strings from TP 8020 without technical changes made by TP 8020. Emphasize that NQNs use UTF-8.
- Clarified that the Transport Requirements field value of 10b (i.e., Not required) in the Discovery Log Page Entry data structure includes both "supported but not required" and "not required due to lack of support".
- Clarified in the FID Supported and Effects data structure that it is the feature identified by the FID that is modified.

## NVM Express NVM Command Set Specification 1.0c:

### Editorial Changes:

- Added the Capacity Exceeded status code to the list of generic command status code values.
- Clarified that PRINFO field in Compare command applies to both the read and the write portions of that command.
- Corrected the referenced section for the Invalid Protection Information status associated with the EILBRT or EILBRT fields.
- Clarified the comparison of the specified NSID as opposed to the specified namespace when comparing to an allocated NSID, an unallocated NSID, or an invalid NSID.

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

## NVM Express Management Interface Specification 1.2c:

### New Requirement:

- Added the explicit requirement that a Responder shall not perform any text processing that is specific to the character set or locale such as checks for byte values not used by UTF-8, Unicode normalization, etc. This requirement was previously implied.

### Editorial Changes:

- Fixed issue introduced by another ECN where CIAP and MEB error checking was incorrect for Response Messages.
- Clarify that the NVMe-MI Send command Completion Queue Entry Dword 0 definition is bits not bytes.
- Clarify that the NVMe-MI Receive command Completion Queue Entry Dword 0 definition is bits not bytes.
- Use a non-breaking hyphen in “NVMe-MI”.
- Expanded details of UTF-8 support by adding improved description of UTF-8 strings from TP 8020 without technical changes made by TP 8020.
- Qualify some stand-alone occurrences of “packet” with “MCTP” (i.e. packet -> MCTP packet).

## NVM Express TCP Transport Specification 1.0c:

### Editorial Changes:

- Clarified the NVMe/TCP Queue Establishment Sequence figure to include the TCP ACK.

## NVM Express Key Value Command Set Specification 1.0c:

### Editorial Changes:

- Revise definition of NUSE to avoid use of the word “allocated”. This removes any implication of interaction with the deallocation functionality of the Sanitize command.

Note:

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

# Description of NVM Express Base Specification 2.0c changes

<Editor: Change all instances of “NVMe-MI” to use a non-breaking hyphen>

Modify section 1 as shown below:

## 1 Introduction

...

### 1.4 Conventions

...

#### 1.4.2 Numerical Descriptions

...

<Note: TP8020 already added the following text>

Some parameters are defined as an ASCII string. ASCII strings shall contain only code values (i.e., byte values or octet values) 20h through 7Eh. For the string “Copyright”, the character “C” is the first byte, the character “o” is the second byte, etc. ~~The ASCII strings~~ **is are** left justified. ~~and If padding is necessary, then the string~~ shall be padded with spaces (i.e., ASCII character 20h) to the right ~~if necessary unless the string is specified as null-terminated.~~

Some parameters are defined as a UTF-8 string. UTF-8 strings shall contain only byte values (i.e., octet values) 20h through 7Eh, 80h through BFh, and C2h through F4h (refer to sections 1 to 3 of RFC 3629). For the string “Copyright”, the character “C” is the first byte, the character “o” is the second byte, etc. UTF-8 strings are left justified. If padding is necessary, then the string shall be padded with spaces (i.e., ASCII character 20h, Unicode character U+0020) to the right unless the string is specified as null-terminated.

If padding is necessary for a field that contains a null-terminated string then the field should be padded with nulls (i.e., ASCII character 00h, Unicode character U+0000) to the right of the string.

...

### 1.5 Definitions

...

#### 1.5.37 ~~n~~Namespace ID (NSID)

An identifier used by a controller to provide access to a namespace or the name of the field in the SQE that contains the namespace identifier (refer to Figure 87). Refer to section 3.2.1 for the definitions of valid NSID, invalid NSID, active NSID, inactive NSID, allocated NSID, and unallocated NSID.

...

#### 1.5.TBD specified namespace

The namespace that is associated with the value specified by the Namespace Identifier (NSID) field in a command as defined by the Common Command Format (refer to **Figure 87**).

...

### 1.8 References

...

RFC 3629, ~~Alis Technologies~~, F. Yergeau, “UTF-8, a transformation format of ISO 10646”, November 2003. Available from <https://www.ietf.org/rfc.html>.

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

...

**Modify section 2 as shown below:**

## 2 Theory of Operation

...

### 2.3 NVM Storage Model

...

#### 2.3.2 I/O Command Sets

...

A controller may support one or more I/O command sets. ~~and may be attached to N~~namespaces that are associated with the I/O command sets that are supported and enabled on a controller may be attached to that controller ~~associated with different I/O command sets~~. A host issues commands to a namespace and those commands are interpreted based on the I/O command set associated with that namespace.

...

### 2.4 Extended Capabilities Theory

#### 2.4.1 Multi-Path I/O and Namespace Sharing

...

This section provides an overview of multi-path I/O and namespace sharing. Multi-path I/O refers to two or more completely independent paths between a single host and a namespace while namespace sharing refers to the ability for two or more hosts to access a common shared namespace using different NVM Express controllers. Both multi-path I/O and namespace sharing require that the NVM subsystem contain two or more controllers. NVM subsystems that support Multi-Path I/O and Namespace Sharing may also support asymmetric controller behavior (refer to section 2.4.2). Concurrent access to a shared namespace by two or more hosts requires some form of coordination between hosts. The procedure used to coordinate these hosts is outside the scope of this specification.

Figure 16 shows an NVM subsystem that contains a single NVM Express controller implemented over PCI Express and a single PCI Express port. Since this is a single Function PCI Express device, the NVM Express controller shall be associated with PCI Function 0. A controller may support multiple namespaces. The controller in Figure 16 supports two namespaces labeled NS A and NS B. Associated with each controller namespace is a namespace ID, labeled as NSID 1 and NSID 2, that is used by the controller to reference a specific namespace. The namespace ID is distinct from the namespace itself and is the handle a host and controller use to specify a particular namespace in a command. The selection of a controller's namespace IDs is outside the scope of this specification. In this example ~~namespace~~NSID 1 is associated with namespace A and ~~namespace~~NSID 2 is associated with namespace B. Both namespaces are private to the controller and this configuration supports neither multi-path I/O nor namespace sharing.



**Figure 16: NVM Express Controller with Two Namespaces**

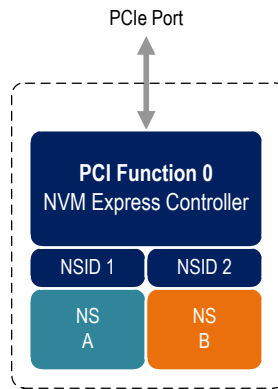
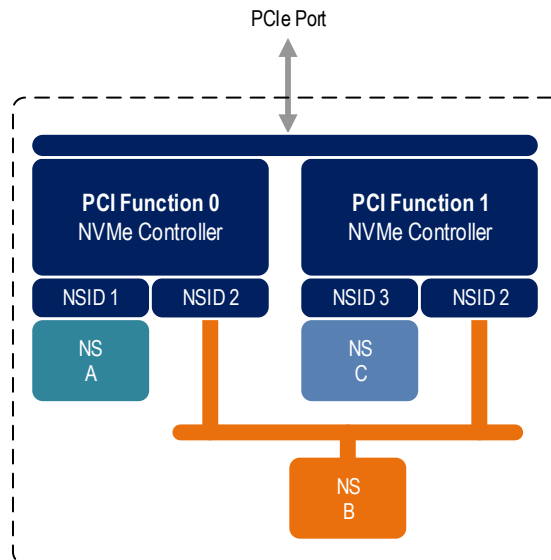


Figure 17 shows a multi-Function NVM subsystem with a single PCI Express port containing two controllers implementing NVMe over PCIe. One controller is associated with PCI Function 0 and the other controller is associated with PCI Function 1. Each controller supports a single private namespace and access to shared namespace B. The namespace ID shall be the same in all controllers that have access to a particular shared namespace. In this example both controllers use ~~namespace~~ NSID 2 to access shared namespace B.

**Figure 17: NVM Subsystem with Two Controllers and One Port**



...

**Modify section 3 as shown below:**

## 3 Admin Command Set

### 3.1 NVM Controller Architecture

...

#### 3.1.1 Controller Model

...

In a dynamic controller model, the controller is allocated by the NVM subsystem on demand. In this model, all controllers allocated to a specific host have the same state at the time the association is

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

established, including ~~attached namespaces and~~ Feature settings. The initial set of attached namespaces should be the same for all controllers that are allocated to a specific host and accessed via the same NVM subsystem port. The initial set of attached namespaces may differ among controllers that are each accessed via a different NVM subsystem port. Changes to a dynamic controller (e.g., attached namespaces, Feature settings) after the association is established do not impact other dynamic controllers.

...

**3.2. NVM Subsystem Entities**

...

**3.2.1 Namespaces**

**3.2.1.6 NSID and Namespace Usage**

...

Namespace IDs may change across power off conditions. However, it is recommended that Namespace IDs remain static across power off conditions in order to avoid issues with host software. To determine if the same namespace has been encountered, the host may use the:

- a) UUID field in the Namespace Identification Descriptor (refer to Figure 277), if present;
- b) NGUID field in the Identify Namespace data (refer to the applicable I/O Command Set specification) or in the Namespace Identification Descriptor, if present; or
- c) EUI64 field in the Identify Namespace data or in the Namespace Identification Descriptor, if present.

...

**3.3 NVM Queue Models**

...

**3.3.2 Message-based Transport Queue Model**


...

**3.3.2.1 Capsules and Data Transfers**

...

**3.3.2.1.1 Command Capsules**

A command capsule is sent from a host to a controller. It contains a submission queue entry (SQE) and may optionally contain data or SGLs. The SQE is 64 bytes in size and contains the Admin command, I/O command, or Fabrics command to be executed.



The Command Identifier field in the SQE shall be unique among all outstanding commands associated with that queue. If there is data or additional SGLs to be transferred within the capsule, then the SGL descriptor in the SQE contains a Data Block, Segment Descriptor, or Last Segment Descriptor specifying

an appropriate Offset address. The definition for the submission queue entry when the command is a Fabrics command is ~~defined shown~~ in Figure 80. The definition for the submission queue entry when the command is an Admin or I/O command is ~~defined shown~~ in section 3.3.3.1, ~~where the Metadata Pointer field is reserved~~. Bytes 03:00 share a common format across commands.

**Figure 80: Fabrics Command Capsule – Submission Queue Entry Format**

Bytes	Description
00	<del>Opcode (OPC): Set to 7Fh to indicate a Fabrics command.</del>
01	<del>PRP or SGL for Data Transfer (PSDT): This field is described in Figure 86. Bits 7:6 of this field should be set to 01b, and may be cleared to 00b.</del>
03:02	<del>Command Identifier (CID): This field specifies a unique identifier for the command. The identifier shall be unique among all outstanding commands associated with a particular queue.</del>
03:00	Command Dword 0: Refer to Figure TBD.
04	<b>Fabrics Command Type (FCTYPE):</b> This field specifies the Fabrics command transferred in the capsule. The Fabrics command types are defined in Figure 375. If this field is set to a reserved value, the command should be aborted with a status code of Invalid Field in Command.
<del>39</del> 23:05	Reserved
39:24	<b>SGL Descriptor 1 (SGL1):</b> This field contains a Transport SGL Data Block descriptor or a Keyed SGL Data Block descriptor that describes the entire data transfer. Refer to section 4.1.2 for the definition of SGL descriptors.  This field is used for Fabrics commands that transfer data. If a Fabrics command does not transfer data, then this field is reserved.
63:40	<b>Fabrics Command Type Specific:</b> This field is Fabrics command type specific.

**Figure TBD: Fabrics Command Capsule – Command Dword 0**

Bits	Description										
31:16	<b>Command Identifier (CID):</b> Refer to <a href="#">Figure 86</a> .										
15:14	<b>PRP or SGL for Data Transfer (PSDT):</b>										
	<table><tr><th>Value</th><th>Description</th></tr><tr><td>00b</td><td>This value is used for Fabrics commands that do not transfer data.  If this value is used for Fabrics commands that transfer data, then SGLs are used for this transfer. A host should use the value 10b rather than 00b for Fabrics commands that transfer data.</td></tr><tr><td>01b</td><td>Reserved</td></tr><tr><td>10b</td><td>This value is used for Fabrics commands that transfer data. SGLs are used for this transfer.</td></tr><tr><td>11b</td><td>Reserved</td></tr></table>	Value	Description	00b	This value is used for Fabrics commands that do not transfer data.  If this value is used for Fabrics commands that transfer data, then SGLs are used for this transfer. A host should use the value 10b rather than 00b for Fabrics commands that transfer data.	01b	Reserved	10b	This value is used for Fabrics commands that transfer data. SGLs are used for this transfer.	11b	Reserved
	Value	Description									
	00b	This value is used for Fabrics commands that do not transfer data.  If this value is used for Fabrics commands that transfer data, then SGLs are used for this transfer. A host should use the value 10b rather than 00b for Fabrics commands that transfer data.									
	01b	Reserved									
	10b	This value is used for Fabrics commands that transfer data. SGLs are used for this transfer.									
	11b	Reserved									
13:10	Reserved										
09:08	<b>Fused Operation (FUSE):</b> Refer to <a href="#">Figure 86</a> . There are no fused Fabrics commands and as a result this field is cleared to 00b.										
07:00	<b>Opcode (OPC):</b> This field is <del>S</del> set to 7Fh to indicate a Fabrics command.										

...

### 3.3.3 Queueing Data Structures

...

#### 3.3.3.2 Common Completion Queue Entry

...

##### 3.3.3.2.1 Status Field Definition

...

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

### 3.3.3.2.1.1 Generic Command Status Definition

...

Figure 94: Status Code – Generic Command Status Values

Value	Description	I/O Command Set Specific	I/O Command Set(s) <sup>1</sup>
...			
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 8.3 8.5.	No	
...			
80h	<b>LBA Out of Range:</b> See the applicable I/O Command Set specification for the description.	Yes	NVM, ZNS
81h	<b>Capacity Exceeded:</b> <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. The command attempted an operation that exceeds the capacity of the namespace.</del>	No	
...			

...

**Modify section 4 as shown below:**

## 4 Data Structures

...

### 4.5 NVMe Qualified Names

...

An NVMe Qualified Name is encoded as a UTF-8 string of Unicode characters (refer to section 1.4.2) with the following properties:

- The encoding is UTF-8 (refer to RFC 3629);
- The following characters are used in formatting:
  - dash ('-'=U+002d);
  - dot ('.'=U+002e); and
  - colon (':'=U+003a);
- The maximum name is 223 bytes in length; and
- The string is null terminated.

...

**Modify section 5 as shown below:**

## 5 Admin Command Set

...

### 5.9 Device Self-test command

...

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

**Figure 171: Device Self-test Namespace Test Action**

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

...

## 5.14 Format NVM command

...

**Figure 188: Format NVM – Operation Scope**

SES	FNA bit 0	FNA bit 1	NSID	Format Operation
000b (i.e., not a secure erase)	0	N/A	FFFFFFFFh <sup>2</sup>	All namespaces attached to the controller. Other namespaces are not affected.
	0		Any allocated value (refer to section 3.2.1.3)	The Particular specified namespace specified. Other namespaces are not affected.
	1 <sup>3</sup>		Any allocated value (refer to section 3.2.1.3) or FFFFFFFFh	All namespaces that exist in the NVM subsystem.
001b or 010b (i.e., secure erase)	N/A	0	FFFFFFFFh <sup>2</sup>	All namespaces attached to the controller. Other namespaces are not affected.
		0	Any allocated value (refer to section 3.2.1.3)	The Particular specified namespace specified. Other namespaces are not affected.
		1 <sup>3</sup>	Any allocated value (refer to section 3.2.1.3) or FFFFFFFFh	All namespaces that exist in the NVM subsystem.
All others				The controller shall abort the command with a status code of Invalid Field in Command
Notes:				
1. For a Format NVM command with Secure Erase, this column refers to bit 1 in the FNA field in the Identify Controller data structure (refer to Figure 275) and bit 0 in the FNA field 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. If bit 3 in the FNA field is set to '1', then this value is not supported.				
3. If bit 3 in the FNA field is set to '1', then this value does not occur. Refer to Figure 275.				

...

## 5.16 Get Log Page command

...

### 5.16.1 Log Specific Information

...

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

### 5.16.1.3 SMART / Health Information (Log Identifier 02h)

...

**Figure 207: SMART / Health Information Log Page**

Bytes	Description
...	
159:144	<p><b>Unsafe Shutdowns:</b> Contains the number of unsafe shutdowns. This count is incremented when the controller does not report it is safe to power down prior to loss of main power.</p> <p>If CAP.CPS is cleared to 00b or set to 01b, it is safe to power down the controller when <del>a-the controller is shutdown processing is complete</del> (i.e., <del>CSTS.ST is cleared to '0' and CSTS.SHST is set to 10b</del>).</p> <p>If CAP.CPS is set to 10b, it is safe to power down the domain when <del>all controllers in that domain are shutdown (e.g., NVM Subsystem Shutdown processing is complete—(i.e., CSTS.ST is set to '1' and CSTS.SHST is set to 10b)</del>.</p> <p>If CAP.CPS is set to 11b, it is safe to power down the NVM subsystem when <del>all controllers in the NVM subsystem are shutdown (e.g., NVM Subsystem Shutdown processing is complete (i.e., CSTS.ST is set to '1' and CSTS.SHST is set to 10b)</del>.</p>
...	

...

### 5.16.1.14 Persistent Event (Log Identifier 0Dh)

...

**Figure 223: Persistent Event Log Specific Parameter Field**

Bits	Description						
14:10	Reserved						
09:08	<p><b>Action:</b> This field specifies the action the controller shall take during processing this Get Log Page command.</p> <table> <tr> <th>Value</th><th>Definition</th></tr> <tr> <td>...</td><td></td></tr> <tr> <td>11b</td><td> <p><b>Establish Context and Read 512 Bytes of Header:</b> The controller shall:</p> <ol style="list-style-type: none"> <li>determine the length of the Persistent Event log page data;</li> <li>determine the set of events to report in the Persistent Event log page data; and</li> <li>if a reporting context does not already exist, then establish a persistent event log reporting context to store information describing the Persistent Event log page data to be reported and track Persistent Event log page data accesses.</li> </ol> <p>If a persistent event log reporting context did not already exist when the Get Log Page command was processed, then the controller shall:</p> <ol style="list-style-type: none"> <li>establish a persistent event log reporting context; and</li> <li>after establishing the context, return 512 bytes of the Persistent Event Log Header starting at offset 0h with the Reporting Context Exists bit cleared to '0' with a status code of Successful Completion.</li> </ol> <p>If a persistent event log reporting context already existed when the Get Log Page command was processed, then the controller shall return 512 bytes of the Persistent Event Log Header starting at offset 0h with the Reporting Context Exists bit set to '1' with a status code of Successful Completion.</p> <p>The 512 bytes of the Persistent Event Log Header shall be returned regardless of the values in the LPOL, LPOU, NUMDL, and NUMDU fields (i.e., the controller shall ignore the LPOL, LPOU, NUMDL, and NUMDU fields). <b>Hosts should ensure allocation of at least 512 bytes in the data buffer to avoid possible data corruption.</b></p> </td></tr> </table>	Value	Definition	...		11b	<p><b>Establish Context and Read 512 Bytes of Header:</b> The controller shall:</p> <ol style="list-style-type: none"> <li>determine the length of the Persistent Event log page data;</li> <li>determine the set of events to report in the Persistent Event log page data; and</li> <li>if a reporting context does not already exist, then establish a persistent event log reporting context to store information describing the Persistent Event log page data to be reported and track Persistent Event log page data accesses.</li> </ol> <p>If a persistent event log reporting context did not already exist when the Get Log Page command was processed, then the controller shall:</p> <ol style="list-style-type: none"> <li>establish a persistent event log reporting context; and</li> <li>after establishing the context, return 512 bytes of the Persistent Event Log Header starting at offset 0h with the Reporting Context Exists bit cleared to '0' with a status code of Successful Completion.</li> </ol> <p>If a persistent event log reporting context already existed when the Get Log Page command was processed, then the controller shall return 512 bytes of the Persistent Event Log Header starting at offset 0h with the Reporting Context Exists bit set to '1' with a status code of Successful Completion.</p> <p>The 512 bytes of the Persistent Event Log Header shall be returned regardless of the values in the LPOL, LPOU, NUMDL, and NUMDU fields (i.e., the controller shall ignore the LPOL, LPOU, NUMDL, and NUMDU fields). <b>Hosts should ensure allocation of at least 512 bytes in the data buffer to avoid possible data corruption.</b></p>
Value	Definition						
...							
11b	<p><b>Establish Context and Read 512 Bytes of Header:</b> The controller shall:</p> <ol style="list-style-type: none"> <li>determine the length of the Persistent Event log page data;</li> <li>determine the set of events to report in the Persistent Event log page data; and</li> <li>if a reporting context does not already exist, then establish a persistent event log reporting context to store information describing the Persistent Event log page data to be reported and track Persistent Event log page data accesses.</li> </ol> <p>If a persistent event log reporting context did not already exist when the Get Log Page command was processed, then the controller shall:</p> <ol style="list-style-type: none"> <li>establish a persistent event log reporting context; and</li> <li>after establishing the context, return 512 bytes of the Persistent Event Log Header starting at offset 0h with the Reporting Context Exists bit cleared to '0' with a status code of Successful Completion.</li> </ol> <p>If a persistent event log reporting context already existed when the Get Log Page command was processed, then the controller shall return 512 bytes of the Persistent Event Log Header starting at offset 0h with the Reporting Context Exists bit set to '1' with a status code of Successful Completion.</p> <p>The 512 bytes of the Persistent Event Log Header shall be returned regardless of the values in the LPOL, LPOU, NUMDL, and NUMDU fields (i.e., the controller shall ignore the LPOL, LPOU, NUMDL, and NUMDU fields). <b>Hosts should ensure allocation of at least 512 bytes in the data buffer to avoid possible data corruption.</b></p>						

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

...

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

...

**Figure 256: FID Supported and Effects Data Structure**

Bits	Description															
31:20	<b>FID Scope (FSP):</b> This field defines the scope for the associated feature identifier. If the value of this field is 0h, then no scope is reported. If this field is non-zero, then only one bit shall be set to '1'.															
	<table> <tr> <th>Bits</th><th>Description</th></tr> <tr> <td>11:6</td><td>Reserved</td></tr> <tr> <td>5</td><td><b>NVM Subsystem Scope:</b> If set to '1', then modifying the <del>attributes value</del> of the <del>feature identified by this</del> FID may impact the whole NVM subsystem. If cleared to '0' and the FSP field is non-zero, then modifying the <del>attributes value</del> of the <del>feature identified by this</del> FID does not impact the whole NVM subsystem.</td></tr> <tr> <td>4</td><td><b>Domain Scope:</b> If set to '1', then modifying the <del>attributes value</del> of the <del>feature identified by this</del> FID may impact a single Domain. If cleared to '0' and the FSP field is non-zero, then modifying the <del>attributes value</del> of the <del>feature identified by this</del> FID does not impact a single Domain.</td></tr> <tr> <td>3</td><td><b>Endurance Group Scope:</b> If set to '1', then modifying the <del>attributes value</del> of the <del>feature identified by this</del> FID may impact Endurance Groups. If cleared to '0' and the FSP field is non-zero, then modifying the <del>attributes value</del> of the <del>feature identified by this</del> FID does not impact Endurance Groups.</td></tr> <tr> <td>2</td><td><b>NVM Set Scope:</b> If set to '1', then modifying the <del>attributes value</del> of the <del>feature identified by this</del> FID may impact NVM Sets. If cleared to '0' and the FSP field is non-zero, then modifying the <del>attributes value</del> of the <del>feature identified by this</del> FID does not impact NVM Sets.</td></tr> <tr> <td>1</td><td><b>Controller Scope:</b> If set to '1', then modifying the <del>attributes value</del> of the <del>feature identified by this</del> FID may impact the controller. If cleared to '0' and the FSP field is non-zero, then the <del>feature identified by this</del> FID does not have controller scope.</td></tr> <tr> <td>0</td><td><b>Namespace Scope:</b> If set to '1', then modifying the <del>attributes value</del> of the <del>feature identified by this</del> FID may impact namespaces. If cleared to '0' and the FSP field is non-zero, then modifying the <del>attributes value</del> of the <del>feature identified by this</del> FID does not impact namespaces.</td></tr> </table>	Bits	Description	11:6	Reserved	5	<b>NVM Subsystem Scope:</b> If set to '1', then modifying the <del>attributes value</del> of the <del>feature identified by this</del> FID may impact the whole NVM subsystem. If cleared to '0' and the FSP field is non-zero, then modifying the <del>attributes value</del> of the <del>feature identified by this</del> FID does not impact the whole NVM subsystem.	4	<b>Domain Scope:</b> If set to '1', then modifying the <del>attributes value</del> of the <del>feature identified by this</del> FID may impact a single Domain. If cleared to '0' and the FSP field is non-zero, then modifying the <del>attributes value</del> of the <del>feature identified by this</del> FID does not impact a single Domain.	3	<b>Endurance Group Scope:</b> If set to '1', then modifying the <del>attributes value</del> of the <del>feature identified by this</del> FID may impact Endurance Groups. If cleared to '0' and the FSP field is non-zero, then modifying the <del>attributes value</del> of the <del>feature identified by this</del> FID does not impact Endurance Groups.	2	<b>NVM Set Scope:</b> If set to '1', then modifying the <del>attributes value</del> of the <del>feature identified by this</del> FID may impact NVM Sets. If cleared to '0' and the FSP field is non-zero, then modifying the <del>attributes value</del> of the <del>feature identified by this</del> FID does not impact NVM Sets.	1	<b>Controller Scope:</b> If set to '1', then modifying the <del>attributes value</del> of the <del>feature identified by this</del> FID may impact the controller. If cleared to '0' and the FSP field is non-zero, then the <del>feature identified by this</del> FID does not have controller scope.	0
Bits	Description															
11:6	Reserved															
5	<b>NVM Subsystem Scope:</b> If set to '1', then modifying the <del>attributes value</del> of the <del>feature identified by this</del> FID may impact the whole NVM subsystem. If cleared to '0' and the FSP field is non-zero, then modifying the <del>attributes value</del> of the <del>feature identified by this</del> FID does not impact the whole NVM subsystem.															
4	<b>Domain Scope:</b> If set to '1', then modifying the <del>attributes value</del> of the <del>feature identified by this</del> FID may impact a single Domain. If cleared to '0' and the FSP field is non-zero, then modifying the <del>attributes value</del> of the <del>feature identified by this</del> FID does not impact a single Domain.															
3	<b>Endurance Group Scope:</b> If set to '1', then modifying the <del>attributes value</del> of the <del>feature identified by this</del> FID may impact Endurance Groups. If cleared to '0' and the FSP field is non-zero, then modifying the <del>attributes value</del> of the <del>feature identified by this</del> FID does not impact Endurance Groups.															
2	<b>NVM Set Scope:</b> If set to '1', then modifying the <del>attributes value</del> of the <del>feature identified by this</del> FID may impact NVM Sets. If cleared to '0' and the FSP field is non-zero, then modifying the <del>attributes value</del> of the <del>feature identified by this</del> FID does not impact NVM Sets.															
1	<b>Controller Scope:</b> If set to '1', then modifying the <del>attributes value</del> of the <del>feature identified by this</del> FID may impact the controller. If cleared to '0' and the FSP field is non-zero, then the <del>feature identified by this</del> FID does not have controller scope.															
0	<b>Namespace Scope:</b> If set to '1', then modifying the <del>attributes value</del> of the <del>feature identified by this</del> FID may impact namespaces. If cleared to '0' and the FSP field is non-zero, then modifying the <del>attributes value</del> of the <del>feature identified by this</del> FID does not impact namespaces.															
19	<b>UUID Selection Supported:</b> If set to '1', then the controller supports the selection of a UUID (refer to section 8.25) by a Get Features command or a Set Features command using this FID. If cleared to '0', then the controller does not support the selection of a UUID by a Get Features command or a Set Features command using this FID.															
18:05	Reserved															
04	<b>Controller Capability Change (CCC):</b> If this bit is set to '1', then changing the <del>attributes value</del> of <del>the feature identified by this</del> FID may change controller capabilities. If this bit is cleared to '0', then changing the <del>attributes value</del> of <del>the feature identified by this</del> FID does not modify controller capabilities. Controller capability changes include a firmware update that changes the capabilities reported in the CAP property.															
03	<b>Namespace Inventory Change (NIC):</b> If this bit is set to '1', then changing the <del>attributes value</del> of <del>the feature identified by this</del> FID may change the number of namespaces or capabilities for multiple namespaces. If this bit is cleared to '0', then changing the <del>attributes value</del> of <del>the feature identified by this</del> FID does not modify the number of namespaces or capabilities for multiple namespaces. Namespace inventory changes include adding or removing namespaces.															
02	<b>Namespace Capability Change (NCC):</b> If this bit is set to '1', then changing the <del>attributes value</del> of <del>the feature identified by this</del> FID may change the capabilities of a single namespace. If this bit is cleared to '0', then changing the <del>attributes value</del> of <del>the feature identified by this</del> FID does not modify any namespace capabilities for the specified namespace. Namespace capability changes include a logical format change.															

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

Figure 256: FID Supported and Effects Data Structure

Bits	Description
01	<b>User Data Content Change (UDCC):</b> If this bit is set to '1', then changing the <del>attributes value</del> of the feature identified by this FID may modify user data content in one or more namespaces. If this bit is cleared to '0', then changing the <del>attributes value</del> of the feature identified by this FID does not modify user data content in any namespace.
00	<b>FID Supported (FSUPP):</b> If this bit is set to '1', then this FID is supported by the controller. If this bit is cleared to '0', then this FID is not supported by the controller and all other fields in this structure shall be cleared to 0h.  Refer to section 3.1.2 for the FID support requirements for each controller type.

...

#### 5.16.1.23 Discovery Log Page (Log Identifier 70h)

...

Figure 264: Discovery Log Page Entry Data Structure

Bytes	Description										
...											
03	<b>Transport Requirements (TREQ):</b> Indicates requirements for the NVMe Transport.  Bits 7:3 are reserved.  Bit 2 if set to '1' indicates that the controller is capable of disabling SQ flow control. A controller that is capable of disabling SQ flow control may accept or reject a host request to disable SQ flow control. If cleared to '0', then the controller requires use of SQ flow control.  Bits 1:0 indicate whether connections shall be made over a fabric secure channel (which includes authentication) (refer to section 8.13). <table> <tr> <th>Value</th><th>Definition</th></tr> <tr> <td>00b</td><td>Not specified</td></tr> <tr> <td>01b</td><td>Required</td></tr> <tr> <td>10b</td><td>Not required (i.e., not supported or supported but not required)</td></tr> <tr> <td>11b</td><td>Reserved</td></tr> </table>	Value	Definition	00b	Not specified	01b	Required	10b	Not required (i.e., not supported or supported but not required)	11b	Reserved
Value	Definition										
00b	Not specified										
01b	Required										
10b	Not required (i.e., not supported or supported but not required)										
11b	Reserved										
...											

...

### 5.17 Identify command

...

#### 5.17.2 Identify Data Structures

...

##### 5.17.2.1 Identify Controller Data Structure (CNS 01h)

...

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

Bytes	I/O <sup>1</sup>	Admin <sup>1</sup>	Disc <sup>1</sup>	Description
...				



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

Bytes	I/O <sup>1</sup>	Admin <sup>1</sup>	Disc <sup>1</sup>	Description								
99:96	M	M	RO <sup>3</sup>	<b>Controller Attributes (CTRATT):</b> This field indicates attributes of the controller.								
				<table><tr><th>Bits</th><th>Description</th></tr><tr><td>...</td><td></td></tr><tr><td>6</td><td><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 3.9.2). 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 3.9.1).</td></tr><tr><td>...</td><td></td></tr></table>	Bits	Description	...		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 3.9.2). 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 3.9.1).	...	
				Bits	Description							
				...								
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 3.9.2). 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 3.9.1).											
...												
...												
321:320	M	M	<div><div></div><div>&lt;From ECN117&gt;</div></div>	<b>Keep Alive Support (KAS):</b> This field indicates the granularity of the Keep Alive Timer in 100 millisecond units (refer to section 3.9). 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 3.9.  < ECN117 changed this definition, but this ECN is not showing those changes as the definition is not being modified by this ECN.>								
...												
524	M	M	R	<b>Format NVM Attributes (FNA):</b> This field indicates attributes for the Format NVM command.  Bits 7:4 are reserved.  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.  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.  Bit 1 indicates whether secure erase functionality applies to all namespaces in the 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 specified namespace-specified. If bit 3 is set to '1', then this bit shall be cleared to '0'.  Bit 0 indicates whether the format operation (excluding secure erase) applies to all namespaces in the NVM subsystem or is specific to a particular namespace. If set to '1', then all namespaces in the 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 the 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'.								
....												
Notes: 1. O/M/R definition: O = Optional, M = Mandatory, R = Reserved. 2. Mandatory for I/O controllers using a message-based transport. Reserved for controllers using a memory-based transport. 3. TBKAS bit is optional for Discovery controllers, and all other bits are reserved for Discovery controllers.												

...

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

A list of Namespace Identification Descriptor structures (refer to Figure 277) is returned to the host for the specified namespace if the specified value in the Namespace Identifier (NSID) field ~~if it is~~ is an active NSID. Namespace Identification Descriptor structures consist of one or more Namespace Identifiers (NID) of various types as indicated by the Namespace Identifier Type (NIDT) field in each descriptor. Each NID is assigned to a namespace at namespace creation and remains fixed throughout the life of that namespace. If the NSID field does not specify an active NSID, then refer to section 3.2.1.5 for the status code to return.

...

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

An I/O Command Set specific Identify Namespace data structure (refer to the applicable I/O Command Set specification) is returned to the host for the specified namespace if the specified value in the Namespace Identifier (NSID) field ~~if the NSID is~~ is an active NSID. If the value in the specified NSID field ~~namespace is~~ specifies 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 274). If the I/O Command Set associated with the specified 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.

...

### 5.17.2.8 I/O Command Set Independent Identify Namespace data structure (CNS 08h)

If the Namespace Identifier (NSID) field specifies an active NSID, then the I/O Command Set Independent Identify Namespace data structure (refer to Figure 280) is returned to the host for that specified namespace. If ~~that the value in the specified NSID field specifies~~ ~~namespace is~~ an inactive NSID, then the controller returns a zero filled data structure.

...

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

The Identify Namespace data structure (refer to the NVM Command Set Specification) is returned to the host for the specified namespace if the specified value in the Namespace Identifier (NSID) field ~~if it is~~ is an allocated NSID. If the value in the specified NSID field ~~namespace is~~ specifies an unallocated NSID, then the controller returns a zero filled data structure. If the specified namespace is not associated with an I/O Command Set that specifies logical blocks (e.g., the NVM Command Set), then the controller shall abort the command with a status code of Invalid I/O Command Set.

If the value in the specified NSID field ~~namespace is~~ specifies an invalid NSID, then the controller shall abort the command with a status code of Invalid Namespace or Format. If the NSID field is set to FFFFFFFFh, then the controller should abort the command with a status code of Invalid Namespace or Format.

...

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

A Controller List (refer to section 4.4.1) 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 of controllers that ~~have are attached to the~~ specified namespace ~~attached specified in the Namespace Identifier (NSID) field~~. If the Namespace Identifier (NSID) field is set to FFFFFFFFh, then the controller should abort the command with a status code of Invalid Field in Command.

...

5.17.2.12 Controller list (CNS 13h)

A Controller List (refer to section 4.4.1) of up to 2,047 controller identifiers of I/O controllers is returned containing a controller identifiers greater than or equal to the value specified in the Controller Identifier (CDW10.CNTID) field. ~~The list contains controller identifiers of controllers in the NVM subsystem that are capable of being attached to namespace(s).~~

...

5.17.2.20 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 section 5.17.2.5) is returned to the host for the specified namespace if the specified value in the Namespace Identifier (NSID) field ~~if it is~~ is an allocated NSID. If the value in the specified NSID field 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 274). If the I/O Command Set associated with the specified namespace ~~specified by the NSID field~~ does not support the specific 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 value in the specified NSID field namespace is an invalid NSID, then the controller shall abort the command with a status code of Invalid Namespace or Format. If the NSID field is set to FFFFFFFFh, then the controller should abort the command with a status code of Invalid Namespace or Format.

...

5.22 Namespace Attachment command

...

5.22.1 Command Completion

...

Figure 296: Namespace Attachment – Command Specific Status Values

Value	Description
18h	Namespace Already Attached: The controller is already attached to the specified namespace <del>specified</del> .
...	

...

Modify section 6 as shown below:

6 Fabrics Command Set

...

6.1 Authentication Receive Command and Response

...

<Note to editor: Note the missing colon being added after the figure number in the figure heading.>

Figure 376: Authentication Receive Command – Submission Queue Entry

Bytes	Description
00	<del>Opcode (OPC): Set to 7Fh to indicate a Fabrics command.</del>
04	<del>Reserved</del>

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

**Figure 376: Authentication Receive Command – Submission Queue Entry**

Bytes	Description
03:02	<del>Command Identifier (CID): This field specifies a unique identifier for the command. Refer to the definition in Figure 80.</del>
03:00	Command Dword 0: Refer to Figure TBD.
04	Fabrics Command Type (FCTYPE): Set to 06h to indicate an Authentication Receive command.
23:05	Reserved
39:24	<del>SGL Descriptor 1 (SGL1): This field contains a Transport SGL Data Block descriptor or a Keyed SGL Data Block descriptor that describes the entire data transfer. Refer to section 4.1.2 for the definition of SGL descriptors. Refer to Figure 80.</del>
40	Reserved
...	

...

## 6.2 Authentication Send Command and Response

...

**Figure 378: Authentication Send Command – Submission Queue Entry**

Bytes	Description
00	<del>Opcode (OPC): Set to 7Fh to indicate a Fabrics command.</del>
04	<del>Reserved</del>
03:02	<del>Command Identifier (CID): This field specifies a unique identifier for the command. Refer to the definition in Figure 80.</del>
03:00	Command Dword 0: Refer to Figure TBD.
04	Fabrics Command Type (FCTYPE): Set to 05h to indicate an Authentication Send command.
23:05	Reserved
39:24	<del>SGL Descriptor 1 (SGL1): This field contains a Transport SGL Data Block descriptor or a Keyed SGL Data Block descriptor that describes the entire data transfer. Refer to section 4.1.2 for the definition of SGL descriptors. Refer to Figure 80.</del>
40	Reserved
...	

...

## 6.3 Connect Command and Response

...

**Figure 380: Connect Command – Submission Queue Entry**

Bytes	Description
00	<del>Opcode (OPC): Set to 7Fh to indicate a Fabrics command.</del>
04	<del>Reserved</del>
03:02	<del>Command Identifier (CID): This field specifies a unique identifier for the command. Refer to the definition in Figure 80.</del>
03:00	Command Dword 0: Refer to Figure TBD.
04	Fabrics Command Type (FCTYPE): Set to 01h to indicate a Connect command.
23:05	Reserved
39:24	<del>SGL Descriptor 1 (SGL1): This field contains a Transport SGL Data Block descriptor or a Keyed SGL Data Block descriptor that describes the entire data transfer. Refer to section 4.1.2 for the definition of SGL descriptors. Refer to Figure 80.</del>

Figure 380: Connect Command – Submission Queue Entry

Bytes	Description
41:40	<b>Record Format (RECFMT):</b> Specifies the format of the Connect command capsule. The format of the record specified in this definition shall be 0h. If the NVM subsystem does not support the value specified, then a status code of Incompatible Format shall be returned.
...	

...

## 6.4 Disconnect Command and Response

...

Figure 384: Disconnect Command ~~and Response~~ – Submission Queue Entry

Bytes	Description
00	<del>Opcode (OPC): Set to 7Fh to indicate a Fabrics command.</del>
01	<del>Reserved</del>
03:02	<del>Command Identifier (CID): This field specifies a unique identifier for the command. Refer to the definition in Figure 80.</del>
03:00	Command Dword 0: Refer to Figure TBD. Byte 01 is cleared to 0h.
04	<b>Fabrics Command Type (FCTYPE):</b> Set to 08h to indicate a Disconnect command.
<del>2339</del> :05	<del>Reserved</del>
<del>39:24</del>	<del>SGL Descriptor 1 (SGL1): This field is reserved, as there is no data transferred by this command.</del>
41:40	<b>Record Format (RECFMT):</b> Specifies the format of the Disconnect command capsule. The format of the record specified in this definition shall be 0h. If the NVM subsystem does not support the value specified, then a status code of Incompatible Format shall be returned.
63:48	Reserved

...

## 6.5 Property Get Command and Response

...

Figure 386: Property Get Command – Submission Queue Entry

Bytes	Description
00	<del>Opcode (OPC): Set to 7Fh to indicate a Fabrics command.</del>
01	<del>Reserved</del>
03:02	<del>Command Identifier (CID): This field specifies a unique identifier for the command. Refer to the definition in Figure 80.</del>
03:00	Command Dword 0: Refer to Figure TBD. Byte 01 is cleared to 0h.
04	<b>Fabrics Command Type (FCTYPE):</b> Set to 04h to indicate a Property Get command.
39:05	Reserved
...	

...

## 6.6 Property Set Command and Response

...

Figure 388: Property Set Command – Submission Queue Entry

Bytes	Description
00	<del>Opcode (OPC): Set to 7Fh to indicate a Fabrics command.</del>
01	<del>Reserved</del>

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

Figure 388: Property Set Command – Submission Queue Entry

Bytes	Description
03:02	<del>Command Identifier (CID): This field specifies a unique identifier for the command. Refer to the definition in Figure 80.</del>
03:00	Command Dword 0: Refer to Figure TBD. Byte 01 is cleared to 0h.
04	Fabrics Command Type (FCTYPE): Cleared to 00h to indicate a Property Set command.
39:05	Reserved
...	

...

**Modify section 8 as shown below:**

## 8 Extended Capabilities

...

### 8.1 Asymmetric Namespace Access Reporting

...

#### 8.1.3 Asymmetric Namespace Access states

...

##### 8.1.3.3 ANA Inaccessible state

...

A controller shall abort commands, other than those described in section 8.1.4, with a status code of Asymmetric Access Inaccessible if those commands are submitted while the relationship between the **specified** namespace ~~specified by the command~~ and the controller processing the command is in this state.

...

##### 8.1.3.4 ANA Persistent Loss state

...

A controller shall abort commands, other than those described in section 8.1.4, with a status code of Asymmetric Access Persistent Loss if those commands are submitted while the relationship between the **specified** namespace ~~specified by the command~~ and the controller processing the command is in this state.

...

##### 8.1.3.5 ANA Change state

...

A controller shall abort commands, other than those described in 8.1.4, with a status code of Asymmetric Access Transition if those commands are submitted while the relationship between the **specified** namespace ~~specified by the command~~ and the controller processing the command is in this state.

...

### 8.7 Directives

...

#### 8.7.3 Streams (Directive Type 01h, Optional)

...

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

8.7.3.1 Directive Receive

...

8.7.3.1.2 Get Status (Directive Operation 02h)

...

The data structure returned is defined in Figure 426. All fields are specific to the specified namespace specified if the NSID value was not set to FFFFFFFFh.

...

8.7.3.1.3 Allocate Resources (Directive Operation 03h)

...

Figure 427: Allocate Resources – Command Dword 12

Bits	Description
31:16	Reserved
15:00	<b>Namespace Streams Requested (NSR):</b> This field specifies the number of stream resources the host is requesting be allocated for exclusive use by the specified namespace specified.

...

Figure 428: Allocate Resources – Completion Queue Entry Dword 0

Bits	Description
31:16	Reserved
15:00	<b>Namespace Streams Allocated (NSA):</b> This field indicates the number of streams resources that have been allocated for exclusive use by the specified namespace specified. The allocated resources are available to all controllers associated with that host.

...

8.19 Reservations

...

Controllers that make up an NVM subsystem shall all have the same support for reservations. Although strongly encouraged, namespaces that make up an NVM subsystem are not all required to have the same support for reservations. For example, some namespaces within a single controller may support reservations while others do not, or the supported reservation types may differ among namespaces. If a controller supports reservations, then the controller shall:

- Indicate support for reservations by returning a '1' in bit 5 of the Optional NVM Command Support (ONCS) field in the Identify Controller data structure;
- Support the Reservation Report command (refer to section 7.5), Reservation Register command (refer to section 7.3), Reservation Acquire command (refer to section 7.27.4), and Reservation Release command (refer to section 7.4);
- Support the Reservation Notification log page;
- Support the Reservation Log Page Available asynchronous events;
- Support the Reservation Notification Mask Feature;
- Support the Host Identifier Feature; and
- Support the Reservation Persistence Feature.

...

8.19.7 Preempting a Reservation or Registration

...

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

A host may abort commands as a side effect of preempting a reservation by executing a Reservation Acquire command (refer to section 7.1) and setting the RACQA field to 010b (Preempt and Abort). The behavior of such a command is exactly the same as that described above with the RACQA field set to 001b (Preempt), with two exceptions:

- After the atomic operation changes namespace reservation and registration state, all controllers associated with any host whose reservation or registration is preempted by that atomic operation are requested to abort all commands being processed that were addressed to the **specified** namespace ~~specified in the Namespace Identifier field (i.e.,~~ the NSID field in the Reservation Acquire command) (refer to section 3.4.4 for the definition of “being processed”); and
- Completion of the Reservation Acquire command shall not occur until all commands that are requested to be aborted are completed, regardless of whether or not each command is actually aborted.

...



# Description of NVM Express NVM Command Set Specification 1.0c changes

Modify section 3 as shown below:

## 3 I/O Commands for the NVM Command Set

...

### 3.1 Submission Queue Entry and Completion Queue Entry

...

#### 3.1.2 NVM Command Set Specific Status Values

...

Figure 15: Status Code – Generic Command Status Values

Value	Description
14h	<b>Atomic Write Unit Exceeded:</b> The length specified exceeds the atomic write unit size.
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 least significant two bits of the Address or Length are not cleared to 00b.  Note: An implementation compliant with revision 1.2.1 of the NVM Express Base Specification or earlier may use the status code value of 15h to indicate SGL Data Block Granularity Invalid.
80h	<b>LBA Out of Range:</b> The command references an LBA that exceeds the size of the namespace.
81h	<b>Capacity Exceeded:</b> The command requested an operation that exceeds the capacity of the namespace. This error occurs when the Namespace Utilization exceeds the Namespace Capacity, as reported in Figure 97.

...

### 3.2 NVM Command Set Commands

...

#### 3.2.1 Compare command

...

Figure 23: Compare – Command Dword 12

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

...

##### 3.2.1.1 Command Completion

...

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

**Figure 26: Compare – Command Specific Status Values**

Value	Description
81h	<b>Invalid Protection Information:</b> The Protection Information (PRINFO) field (refer to Figure 23) settings specified in the command are invalid for the Protection Information with which the namespace was formatted (refer to the PI field in Figure 78 and the DPS field in Figure 97) or the EILBRT field is invalid (refer to section 5.2.3 5.2.2.5).

...

### 3.2.2 Copy command

...

If:

- the Copy Descriptor Format specified in the Descriptor Format field is supported by the controller;
- the specified namespace ~~specified by NSID~~ is formatted to use 16b Guard Protection Information; and
- the Descriptor Format is not cleared to 0h,

then the command shall be aborted with the status code of Invalid Namespace or Format.

If:

- the Copy Descriptor Format specified in the Descriptor Format field is supported by the controller;
- the specified namespace ~~specified by NSID~~ is formatted to use 32b Guard Protection Information or 64b Guard Protection Information; and
- the Descriptor Format is not set to 1h,

then the command shall be aborted with the status code of Invalid Namespace or Format.

...

#### 3.2.2.1 Command Completion

...

**Figure 37: Copy – Command Specific Status Values**

Value	Description
81h	<b>Invalid Protection Information:</b> The protection information specified by the command is invalid due to: <ul style="list-style-type: none"> <li>The Protection Information Read (PRINFO) field or Protection Information Write (PRINFOW) field (refer to Figure 30) containing an invalid value for the Protection Information with which the namespace was formatted (refer to the PI field in the Format NVM Command section in the NVM Express Base Specification and the DPS field in Figure 97)</li> <li>the ILBRT field being invalid (refer to section 5.2.3 5.2.2.5); or</li> <li>the EILBRT field in a Source Range Entry being invalid (refer to section 5.2.3 5.2.2.5).</li> </ul>
...	...

...

### 3.2.4 Read command

...

#### 3.2.4.1 Command Completion

...

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

**Figure 52: Read – Command Specific Status Values**

Value	Description
80h	<b>Conflicting Attributes:</b> The attributes specified in the command are conflicting.
81h	<b>Invalid Protection Information:</b> The Protection Information (PRINFO) field (refer to Figure 48) settings specified in the command are invalid for the Protection Information with which the namespace was formatted (refer to the PI field in Figure 78 and the DPS field in Figure 97) or the EILBRT field is invalid (refer to section 5.2.3 5-2-2-5).

...

### 3.2.5 Verify command

...

#### 3.2.5.1 Command Completion

...

**Figure 58: Verify – Command Specific Status Values**

Value	Description
81h	<b>Invalid Protection Information:</b> The Protection Information (PRINFO) field (refer to Figure 55) settings specified in the command are invalid for the Protection Information with which the namespace was formatted (refer to the PI field in Figure 78 and the DPS field in Figure 97) or the EILBRT field is invalid (refer to section 5.2.3 5-2-2-5).

...

### 3.2.6 Write command

...

#### 3.2.6.1 Command Completion

...

**Figure 67: Write – Command Specific Status Values**

Value	Description
80h	<b>Conflicting Attributes:</b> The attributes specified in the command are conflicting.
81h	<b>Invalid Protection Information:</b> The Protection Information (PRINFO) field (refer to Figure 63) settings specified in the command are invalid for the Protection Information with which the namespace was formatted (refer to the PI field in Figure 78 and the DPS field in Figure 97) or the ILBRT field is invalid (refer to section 5.2.3 5-2-2-5).
...	...

...

### 3.2.8 Write Zeroes command

...

#### 3.2.8.1 Command Completion

...

Figure 76: Write Zeroes – Command Specific Status Values

Value	Description
81h	<b>Invalid Protection Information:</b> The Protection Information (PRINFO) field (refer to Figure 73) settings specified in the command are invalid for the Protection Information with which the namespace was formatted (refer to the PI field in Figure 78 and the DPS field in Figure 97) or the ILBRT field is invalid (refer to section 5.2.3 5.2.2.5).
...	

...

*Modify section 4 as shown below:*

## 4 Admin Commands for the NVM Command Set

...

### 4.1 Admin Command behavior for the NVM Command Set

...

#### 4.1.3 Get Features & Set Features commands

...

##### 4.1.3.2 Error Recovery (Feature Identifier 05h)

...

Figure 84: Error Recovery – Command Dword 11

Bits	Description
31:17	Reserved
16	<b>Deallocated or Unwritten Logical Block Error Enable (DULBE):</b> If set to '1', then the Deallocated or Unwritten Logical Block error is enabled for the specified namespace <del>specified in the NSID field</del> . If cleared to '0', then the Deallocated or Unwritten Logical Block error is disabled for the specified namespace <del>specified in the NSID field</del> . Host software shall only enable this error if the DAE bit in the NSFEAT field is set to '1' in the Identify Namespace data structure. The default value for this bit shall be '0'. Refer to section 3.2.3.2.1.
15:00	<b>Time Limited Error Recovery (TLER):</b> Indicates a limited retry timeout value in 100 millisecond units. This limit applies to I/O commands that support the Limited Retry bit and that are sent to the namespace for which this Feature has been set. The timeout starts when error recovery actions have started while processing the command. A value of 0h indicates that there is no timeout.  Note: This mechanism is primarily intended for use by host software that may have alternate means of recovering the data.

...

#### 4.1.5 Identify Command

...

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

If the Namespace Identifier (NSID) field specifies an active NSID, then the NVM Command Set Identify Namespace data structure (refer to Figure 97) is returned to the host for that specified namespace. If that ~~specified value in the NSID field 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 ~~the~~ status code of Invalid I/O Command Set.

...

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

An Identify Namespace data structure (refer to Figure 97) is returned to the host for the **specified** namespace **if the specified value** in the Namespace Identifier (NSID) field **if it** is an allocated NSID. If the **value in the specified NSID field namespace-is** specifies an unallocated NSID, then the controller returns a zero filled data structure.

If the **value in the specified NSID field namespace-is** specifies an invalid NSID, then the controller shall abort the command with a status code of Invalid Namespace or Format. If the NSID field is set to FFFFFFFFh, then the controller should abort the command with a status code of Invalid Namespace or Format.

...

## 4.2 I/O Command Set Specific Admin commands

...

### 4.2.1 Get LBA Status command

...

If the value in the Action Type (ATYPE) field is set to 10h, then:

- a) the controller shall generate a list of Untracked LBAs as described in this section;
- b) the controller shall return Untracked LBAs and Tracked LBAs in the range specified in the Get LBA Status command for the **specified** namespace ~~specified in the Namespace Identifier (CDW1.NSID);~~
- c) the controller shall remove all LBAs in the range specified in the Get LBA Status command, which prior to processing the Get LBA Status command were successfully re-written, from relevant internal data structures (e.g., internal Tracked LBA list);
- d) the controller shall ensure that any such successfully re-written logical blocks are not reported in any LBA Status Descriptor Entries returned by the Get LBA Status command unless, after having been removed from relevant internal data structures and prior to processing the Get LBA Status command, those LBAs were newly detected as being Potentially Unrecoverable LBAs; and
- e) the list of Untracked LBAs returned by the Get LBA Status command may be discarded by the controller or added to the Tracked LBA list once the command has completed.

If the value in the Action Type (ATYPE) field is set to 11h, then the controller shall:

- a) return Tracked LBAs in the range specified in the Get LBA Status command for the **specified** namespace ~~specified in the Namespace Identifier (CDW1.NSID) field;~~
- b) remove all LBAs in the range specified in the Get LBA Status command, which prior to processing the Get LBA Status command were successfully re-written, from relevant internal data structures (e.g., internal Tracked LBA list);
- c) ensure that any such successfully re-written logical blocks are not reported in any LBA Status Descriptor Entries returned by the Get LBA Status command unless, after having been removed from relevant internal data structures and prior to processing the Get LBA Status command, those LBAs were newly detected as being Potentially Unrecoverable LBAs; and
- d) not perform a foreground scan to generate and return Untracked LBAs.

...

***Modify section 5 as shown below:***

## **5 Extended Capabilities**

...

### **5.2 End-to-end Data Protection**

...

#### **5.2.2 PRACT Bit**

...

##### **5.2.2.4 Protection Information and Compare commands**

...

For the portion of the Compare command that transfers data and protection information from the host to the controller, the protection information checks performed by the controller parallels the Write command protection information checks (refer to section 5.2.2.1). For the portion of the Compare command that transfers data and protection information from the NVM media to the controller, the protection information checks performed by the controller parallels the Read command protection information checks (refer to section 5.2.2.2). The ELBST, EILBRT, PRINFO, STC, ELBATM, and ELBAT fields in the command are used by both sets of protection information checks as defined in section 5.2.3.

...

# Description of NVM Express Management Interface Specification

## 1.2c changes

*<Editor: Change all instances of “NVMe-MI” to use a non-breaking hyphen>*

*Modify section 1 as shown below:*

## 1 Introduction

...

### 1.4 NVM Subsystem Architectural Model

...

Each Management Endpoint advertises the unique set of capabilities supported by that Management Endpoint. All Management Endpoints may support the same commands even though PCIe ports are full duplex with much higher data rates than SMBus ~~(i.e., both SMBus/I2C and PCIe VDM are capable of providing the same functionality).~~

...

### 1.7 Conventions

...

~~Some parameters are defined as a string of ASCII or UTF-8 characters. ASCII strings shall contain only code values 20h to 7Eh. UTF-8 is backwards compatible with ASCII encoding and supports additional characters with variable length encoding. For the string “Copyright”, the character “C” is the least-significant byte, the character “o” is the second byte, etc. The string is left justified and shall be padded with spaces (ASCII character 20h) to the right if necessary.~~

Some parameters are defined as an ASCII string. ASCII strings shall contain only code values (i.e., byte values or octet values) 20h through 7Eh. For the string “Copyright”, the character “C” is the first byte, the character “o” is the second byte, etc. ASCII strings are left justified. If padding is necessary, then the string shall be padded with spaces (i.e., ASCII character 20h) to the right unless the string is specified as null-terminated.

Some parameters are defined as a UTF-8 string. UTF-8 strings shall contain only byte values (i.e., octet values) 20h through 7Eh, 80h through BFh, and C2h through F4h (refer to sections 1 to 3 of RFC 3629). For the string “Copyright”, the character “C” is the first byte, the character “o” is the second byte, etc. UTF-8 strings are left justified. If padding is necessary, then the string shall be padded with spaces (i.e., ASCII character 20h, Unicode character U+0020) to the right unless the string is specified as null-terminated.

If padding is necessary for a field that contains a null-terminated string then the field should be padded with nulls (i.e., ASCII character 00h, Unicode character U+0000) to the right of the string.

For any ASCII string or UTF-8 string received from a Requester, a Responder shall treat that string as a binary string (e.g., it shall not perform any text processing that is specific to the character set or locale such as checks for byte values not used by UTF-8, Unicode normalization, etc.).

...

### 1.11 References

...

PCI-SIG PCI Express SFF-8639 Module Specification, Revision 3.0, Version 1.0. Available from <https://www.pcisig.com>.

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

RFC 3629, F. Yergeau, "UTF-8, a transformation format of ISO 10646", November 2003. Available from <https://www.ietf.org/rfc.html>.

SNIA Native NVMe-oF™ Drive Specification, Version 1.0.1. Available from <https://www.snia.org>.

...

**Modify section 3 as shown below:**

## 3 Message Transport

...

### 3.1 NVMe-MI Messages

...

#### 3.1.1 Message Fields

...

If the Management Endpoint supports the Management Endpoint Buffer (i.e., the Management Endpoint Buffer Size field is set to a non-zero value), then in the out-of-band mechanism, ~~the~~ the Management Endpoint Buffer (MEB) bit in the Message Header specifies whether Message Data is contained in the associated Message Data field of an NVMe-MI Message or in the Management Endpoint Buffer. ~~This bit should only be set to '1' in Command Messages that support Management Endpoint Buffer operation (i.e., those listed in the Management Endpoint Buffer Supported Command List data structure). If the MEB bit is set to '1' in any other Command Message, then the Management Endpoint shall respond with an Invalid Parameter Error Response with the PEL field indicating the MEB bit.~~

The Command Initiated Auto Pause Supported (CIAPS) bit in **Figure 96** indicates if the port supports the Command Initiated Auto Pause (CIAP) bit in Command Messages. If the Command Initiated Auto Pause bit is supported (i.e., the CIAPS bit is set to '1'), then in the out-of-band mechanism, ~~the~~ the Command Initiated Auto Pause (CIAP) bit in the Message Header of a Command Message specifies whether or not the Management Endpoint is automatically paused when a Command Message enters the Process state. A Command Message **in the out-of-band mechanism with the CIAPS bit set to '1' and** with the CIAP bit set to '1' shall be treated by the Management Endpoint as if an implicit Pause Control Primitive, as described in section 4.2.1.1, was received in the Process state with the exception that the Management Endpoint shall not transmit a Control Primitive Response Message. ~~The Command Initiated Auto Pause Supported (CIAPS) bit in Figure 96 indicates if the port supports the Command Initiated Auto Pause (CIAP) bit in Command Messages.~~

...



Figure 19: NVMe-MI Message Fields

Bytes		Description																					
...																							
2	1	Bits	Description																				
		7:2	Reserved																				
		<b>Command Initiated Auto Pause (CIAP):</b> If this bit is set to '1' in a Command Message and the Command Initiated Auto Pause Supported (CIAPS) bit is set to '1', then the Management Endpoint shall be automatically paused when the Command Message enters the Process state. If this bit is cleared to '0' in a Command Message, the Management Endpoint shall not be automatically paused when the Command Message enters the Process state. The usage requirements for this bit are as follows:																					
		<table><tr><th colspan="2">Mechanism</th><th>CIAP Value</th><th>Usage Requirement</th></tr><tr><td rowspan="4">Out-of-band</td><td rowspan="2">Command Messages</td><td>0</td><td>This value is permitted.</td></tr><tr><td>1</td><td>If the CIAPS bit is set to '1', then This a value of '1' for the CIAP bit is permitted.  If the CIAPS bit is cleared to '0', then: a) a value of '1' for the CIAP bit is prohibited; and b) if a value of '1' for the CIAP bit is received in a Command Message, then an Invalid Parameter Error Response with the PEL field indicating this bit should be returned.  &lt;Note to Editor: Convert the "should" in the prior sentence to a "shall" in the non-errata version of the integration spec.&gt;</td></tr><tr><td rowspan="2">Any NVMe-MI Message other than a Command Message</td><td>0</td><td>This value is required.</td></tr><tr><td>1</td><td>This value is prohibited. If this value is received in a Control Primitive, then AAn Invalid Parameter Error Response with the PEL field indicating this bit shall be returned.</td></tr><tr><td rowspan="2">In-band Tunneling</td><td>0</td><td>This value is required.</td></tr><tr><td>1</td><td>This value is prohibited. If this value is received in a Request Message, then AAn Invalid Parameter Error Response with the PEL field indicating this bit shall be returned.</td></tr></table>		Mechanism		CIAP Value	Usage Requirement	Out-of-band	Command Messages	0	This value is permitted.	1	If the CIAPS bit is set to '1', then This a value of '1' for the CIAP bit is permitted.  If the CIAPS bit is cleared to '0', then: a) a value of '1' for the CIAP bit is prohibited; and b) if a value of '1' for the CIAP bit is received in a Command Message, then an Invalid Parameter Error Response with the PEL field indicating this bit should be returned.  <Note to Editor: Convert the "should" in the prior sentence to a "shall" in the non-errata version of the integration spec.>	Any NVMe-MI Message other than a Command Message	0	This value is required.	1	This value is prohibited. If this value is received in a Control Primitive, then AAn Invalid Parameter Error Response with the PEL field indicating this bit shall be returned.	In-band Tunneling	0	This value is required.	1	This value is prohibited. If this value is received in a Request Message, then AAn Invalid Parameter Error Response with the PEL field indicating this bit shall be returned.
		Mechanism		CIAP Value	Usage Requirement																		
		Out-of-band	Command Messages	0	This value is permitted.																		
				1	If the CIAPS bit is set to '1', then This a value of '1' for the CIAP bit is permitted.  If the CIAPS bit is cleared to '0', then: a) a value of '1' for the CIAP bit is prohibited; and b) if a value of '1' for the CIAP bit is received in a Command Message, then an Invalid Parameter Error Response with the PEL field indicating this bit should be returned.  <Note to Editor: Convert the "should" in the prior sentence to a "shall" in the non-errata version of the integration spec.>																		
			Any NVMe-MI Message other than a Command Message	0	This value is required.																		
				1	This value is prohibited. If this value is received in a Control Primitive, then AAn Invalid Parameter Error Response with the PEL field indicating this bit shall be returned.																		
		In-band Tunneling	0	This value is required.																			
1	This value is prohibited. If this value is received in a Request Message, then AAn Invalid Parameter Error Response with the PEL field indicating this bit shall be returned.																						
<b>Management Endpoint Buffer (MEB):</b> This bit indicates whether the Message Data in a Command Message is contained in the Message Data field of this NVMe-MI Message or in the Management Endpoint Buffer. Refer to section 3.1.																							
0	<table><tr><th>Value</th><th>Description</th></tr><tr><td>0b</td><td>The Message Data is contained in the Message Data of this NVMe-MI Message.</td></tr><tr><td>1b</td><td>If the Management Endpoint supports the Management Endpoint Buffer, then The Message Data is contained in the Management Endpoint Buffer.</td></tr></table>		Value	Description	0b	The Message Data is contained in the Message Data of this NVMe-MI Message.	1b	If the Management Endpoint supports the Management Endpoint Buffer, then The Message Data is contained in the Management Endpoint Buffer.															
Value	Description																						
0b	The Message Data is contained in the Message Data of this NVMe-MI Message.																						
1b	If the Management Endpoint supports the Management Endpoint Buffer, then The Message Data is contained in the Management Endpoint Buffer.																						
The usage requirements for this bit are as follows:																							

Bytes	Description																	
	<table border="1"> <thead> <tr> <th>Mechanism</th><th>MEB Value</th><th>Usage Requirement</th></tr> </thead> <tbody> <tr> <td rowspan="4">Out-of-band  Command Messages</td><td>0</td><td>This value is permitted.</td></tr> <tr> <td>1</td><td>If the Management Endpoint supports the Management Endpoint Buffer, then This—a value of ‘1’ for the MEB bit is permitted.  If the Management Endpoint does not support the Management Endpoint Buffer, then:     a) a value of ‘1’ for the MEB bit is prohibited; and     b) if a value of ‘1’ for the MEB bit is received in a Command Message, then an Invalid Parameter Error Response with the PEL field indicating this bit should be returned.  &lt;Note to Editor: Convert the “should” in the prior sentence to a “shall” in the non-errata version of the integration spec.&gt;</td></tr> <tr> <td>0</td><td>This value is required.</td></tr> <tr> <td>1</td><td>This value is prohibited. If this value is received in a Control Primitive, then Aan Invalid Parameter Error Response with the PEL field indicating this bit shall be returned.</td></tr> <tr> <td rowspan="2">In-band Tunneling</td><td>0</td><td>This value is required.</td></tr> <tr> <td>1</td><td>This value is prohibited. If this value is received in a Request Message, then Aaan Invalid Parameter Error Response with the PEL field indicating this bit shall be returned.</td></tr> </tbody> </table>	Mechanism	MEB Value	Usage Requirement	Out-of-band  Command Messages	0	This value is permitted.	1	If the Management Endpoint supports the Management Endpoint Buffer, then This—a value of ‘1’ for the MEB bit is permitted.  If the Management Endpoint does not support the Management Endpoint Buffer, then: a) a value of ‘1’ for the MEB bit is prohibited; and b) if a value of ‘1’ for the MEB bit is received in a Command Message, then an Invalid Parameter Error Response with the PEL field indicating this bit should be returned.  <Note to Editor: Convert the “should” in the prior sentence to a “shall” in the non-errata version of the integration spec.>	0	This value is required.	1	This value is prohibited. If this value is received in a Control Primitive, then Aan Invalid Parameter Error Response with the PEL field indicating this bit shall be returned.	In-band Tunneling	0	This value is required.	1	This value is prohibited. If this value is received in a Request Message, then Aaan Invalid Parameter Error Response with the PEL field indicating this bit shall be returned.
Mechanism	MEB Value	Usage Requirement																
Out-of-band  Command Messages	0	This value is permitted.																
	1	If the Management Endpoint supports the Management Endpoint Buffer, then This—a value of ‘1’ for the MEB bit is permitted.  If the Management Endpoint does not support the Management Endpoint Buffer, then: a) a value of ‘1’ for the MEB bit is prohibited; and b) if a value of ‘1’ for the MEB bit is received in a Command Message, then an Invalid Parameter Error Response with the PEL field indicating this bit should be returned.  <Note to Editor: Convert the “should” in the prior sentence to a “shall” in the non-errata version of the integration spec.>																
	0	This value is required.																
	1	This value is prohibited. If this value is received in a Control Primitive, then Aan Invalid Parameter Error Response with the PEL field indicating this bit shall be returned.																
In-band Tunneling	0	This value is required.																
	1	This value is prohibited. If this value is received in a Request Message, then Aaan Invalid Parameter Error Response with the PEL field indicating this bit shall be returned.																

...

## 4 Message Servicing Model

...

...

...

34

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

...

**Figure 51: NVMe-MI Send – Completion Queue Entry Dword 0 (NSCQED0)**

BytesBits	Description
31:08	<b>Tunneled NVMe Management Response (TNMRESP):</b> This field contains the NVMe Management Response field from the NVMe-MI Command that is being tunneled in-band. If any errors are detected in the NVMe Context as described in section 4.3.1.2, then this field shall be cleared to 0h.
07:00	<b>Tunneled Status (TSTAT):</b> This field contains the Status field from the NVMe-MI Command that is being tunneled in-band. If any errors are detected in the NVMe Context as described in section 4.3.1.2, then this field shall be cleared to 0h.

...

#### 4.3.2 NVMe-MI Receive Command

...

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

...

The definition of Dword 0 of the completion queue entry is in Figure 56.

**Figure 56: NVMe-MI Receive – Completion Queue Entry Dword 0 (NRCQED0)**

BytesBits	Description
31:08	<b>Tunneled NVMe Management Response (TNMRESP):</b> This field contains the NVMe Management Response field from the NVMe-MI Command that is being tunneled in-band. If any errors are detected in the NVMe Context as described in section 4.3.2.2, then this field shall be cleared to 0h.
07:00	<b>Tunneled Status (TSTAT):</b> This field contains the Status field from the NVMe-MI Command that is being tunneled in-band. If any errors are detected in the NVMe Context as described in section 4.3.2.2, then this field shall be cleared to 0h.

...

# Description of NVM Express TCP Transport Specification 1.0c changes

*Modify section 3 as shown below:*

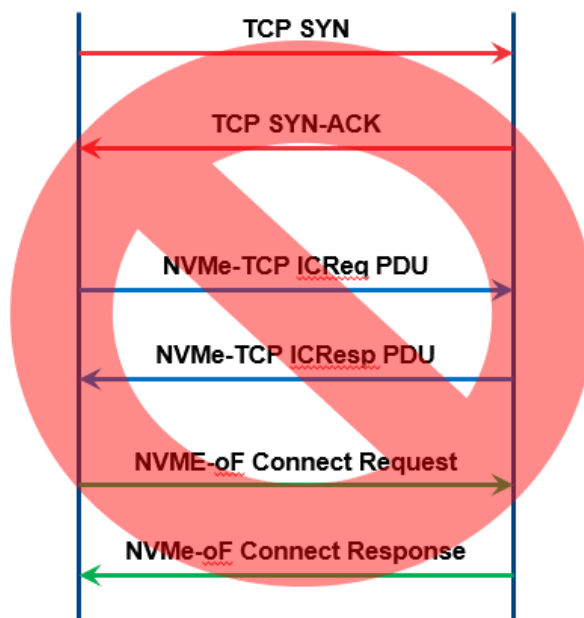
## 3 Transport Binding

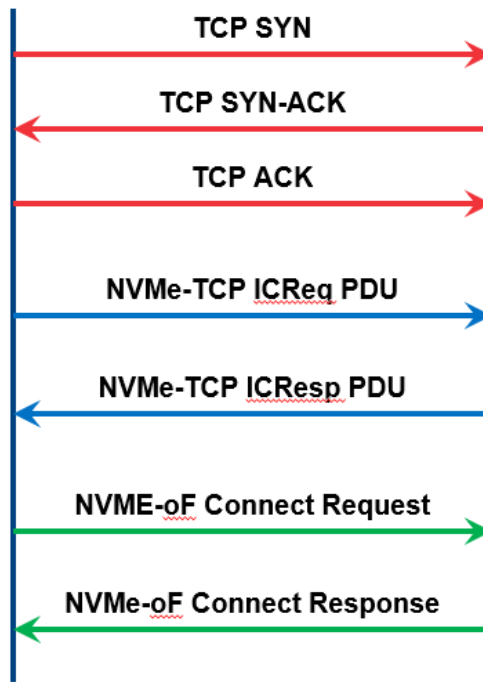
...

### 3.1 Setup & Initialization

...

Figure 7: NVMe/TCP Queue Establishment Sequence





# Description of NVM Express Key Value Command Set Specification

## 1.0c changes

*Modify section 2 as shown below:*

## 2 Key Value Command Set Value

...

### 2.1 Theory of operation

...

#### 2.1.1 Namespaces

...

The number of bytes required to store a given key value pair is greater than or equal to the sum of the size of the KV key and the size of the KV value. Namespace Size and Namespace Utilization reflect the number of bytes required to store the KV value and KV key.

The Key Value Command Set specific Identify Namespace data structure (refer to 4.1.5.1) contains related fields reporting the Namespace Size, Capacity and Utilization:

- The Namespace Size (NSZE) field defines the total size of the namespace in bytes.
- The Namespace Utilization (NUSE) field defines the number of bytes of namespace capacity that are currently allocated in use to store KV keys and KV values in the namespace.

...

*Modify section 4 as shown below:*

## 4 Admin Commands for the Key Value Command Set

### 4.1 Admin Command behavior for the Key Value Command Set

...

#### 4.1.5 Identify Command

...

##### 4.1.5.1 I/O Command Set Specific Identify Namespace Data Structure (CNS 05h, CSI 01h)

...

**Figure 39: Identify – I/O Command Set Specific Identify Namespace Data Structure, Key Value Type Specific**

Bytes	O/M <sup>1</sup>	Description
23:16	M	<p><b>Namespace Utilization (NUSE):</b> This field indicates the current number of bytes <del>allocated of namespace capacity that are in use in the namespace. This is the space</del> to store KV keys and KV values. This field is less than or equal to the Namespace Size field.</p> <p>A key value pair <del>is allocated</del> begins to use namespace capacity when <del>it the</del> key value pair is written with a Store command. A key value pair <del>is deallocated</del> ceases to use namespace capacity when the key value pair is deleted using the Delete command.</p> <p>If the controller supports Asymmetric Namespace Access Reporting (refer to the CMIC field), and the relationship between the controller and the namespace is in the ANA Inaccessible state (refer to the NVM Express Base Specification) or the ANA Persistent Loss state (refer to the NVM Express Base Specification), then this field shall be cleared to 0h.</p>
...	...	...
...	...	...

...