



LEGAL NOTICE:

© **Copyright 2007 - 2021 NVM Express, Inc. ALL RIGHTS RESERVED.**

This erratum to the NVM Express Base Specification revision 2.0 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 to the NVM Express revision 1.4 specification subject, however, to the Member’s continued compliance with the Company’s Intellectual Property Policy and Bylaws and the Member’s Participation Agreement.

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

LEGAL DISCLAIMER:

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

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

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

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

NVM Express™ Technical Errata

Errata ID	101
Revision Date	8/26/2021
Affected Spec Ver.	<p>NVM Express™ Base Specification Revision 2.0</p> <p>NVM Express™ NVM Command Set Specification Revision 1.0</p> <p>NVM Express™ Key Value Command Set Specification Revision 1.0</p> <p>NVM Express™ Zoned Namespace Command Set Specification Revision 1.1</p> <p>NVM Express™ PCI Express® Transport Specification Revision 1.0</p> <p>NVM Express™ RDMA Transport Specification Revision 1.0</p> <p>NVM Express™ TCP Transport Revision 1.0</p> <p>NVM Express™ Management Interface Specification Revision 1.2</p>
Corrected Spec Ver.	

Errata Author(s)

Name	Company
Fred Knight	NetApp
Paul Suhler, John Geldman	Kioxia
Mike Allison	Intel
Judy Brock, Mike Allison	Samsung
Yoni Shternhell	WDC

Errata Overview

This ECN updates and clarifies various text within the NVM Express Base specification, the NVM Express NVM Command Set specification, the NVM Express Zoned Namespace Command Set specification, the NVM Express Key Value Command Set specification, the NVM Express PCI Express Transport specification, the NVM Express RDMA Transport specification, the NVM Express TCP Transport specification, and the NVM Express Management Interface specification.

Revision History

Revision Date	Change Description
5/08/2021	Initial creation
5/11/2021	Add Log Page and Set Feature pointers to Command Set specs (as appropriate)
5/13/2021	Move User Data Write Access definition to future ECN. Define Host Data Read commands and correct SMART data units read/written and host read/write command fields for backward compatibility
5/17/2021	Move "Description of changes" text to new location.
5/19/2021	Adjust name used in the SMART/EG statistics fields. Clean up "User Data Format" term being used with 2 different meanings and no definition (by creating Format Index term). Add "NVM Command Set" to description of Copy Command. Fix FLBAS bits 3:0 description.
5/21/2021	Added in and corrected trademarks and attributions in all eight specifications. Accept changes. Clean up some spacing and pagination, correct some Figure numbers to match the preratification documents. Moved © in -MI to first reference.
5/24/2021	Added 2 commas to each Figure 1 (e.g., ...). Fix intro sentence to Figure 1 so it references the spec it is contained in, rather than referencing the "NVM Command Set" specification. Add "NVM Express" to make the spec name complete, rather than partial. Removed extra © in -MI.
6/2/2021	Updates based on published 2.0 specs + 30-day review feedback (change NSABP to OPTPERF in the description of NPWG, NPWA, NPDG, NPDA, and NOWS fields – to fix a cut/paste error during 2.0 work and restore 1.4 compatibility).
6/8/2021	Add PIP description update (change 1.4 or later to 2.0 or later).
6/9/2021	Correct RMEDIA duplicate bit 0 (delete it).
6/10/2021	Review at TWG, accept tracked changes.
6/18/2021	More 30 day review comments: Update NVM Command Set spec section 2.1.4.2.2 on non-volatility requirements. Fix 2.0 final integration errors in ZNS – write operations term and User Data Out command term.
7/1/2021	Review at TWG, accept tracked changes (except the rewordings in section 2.1.4.2.2 – which were moved to ECN-002).
7/21/2021	Integrated into the NVMe Base Specification Revision 2.0, NVMe NVM Command Set Specification Revision 1.0, NVMe Key Value Command Set Specification Revision 1.0, NVMe Zoned Namespace Command Set Specification Revision 1.1, NVMe PCI Express® Transport Specification Revision 1.0, NVMe RDMA Transport Specification Revision 1.0, NVMe TCP Transport Revision 1.0, and NVMe Management Interface Specification Revision 1.2
7/22/2021	Added the row label in the I/O Command Set specific Identify Namespace data structure in the NVM Command Set Specification.
8/26/2021	Document name updated from ECN-001 to ECN 101 in order to align with NVMe 2.0 nomenclature.

Description of Changes

NVM Express® Base specification:

Added and corrected trademark attributions.

Correct definitions of User Data Read Access command and new Host Data Read command to align SMART and EG log page statistics for backward compatibility and name the list of commands in a less confusing way.

Clarify the distinction between the actual User Data format, and the Format Index used to specify which User Data format to use, rather than use the same term for both concepts.

The ONCS bit 4 set to '1' is defined to indicate support for using the select field for the Get Features command. One portion of the text failed to repeat this requirement; that restatement was added.

The ONCS bit 8 referenced the "Copy command", when it should have referenced the "NVM Command Set Copy command".

Clarify FORMAT NVM command use if the host never sets or clears the LBAFEE field. Associate the already defined Invalid Format error with a host "shall" statement.

Update Log Page value figure (Figure 202) to point to the appropriate command set as necessary for command set specific values.

Telemetry Data is defined as being controller scoped or subsystem scoped. One portion of the text failed to repeat this statement; that restatement was added.

The description of the Offset Index for the ANA Log page was confusing; that section of the text was clarified.

Several entries in the references column of the IDENTIFY CNS value table (Figure 273) were updated to point to the correct sections.

The Capacity Management capability uses the TNVMCAP field and UNVMCAP field (Figure 275). That requirement was restated in the definition of those fields in Identify Controller (CNS=01h).

Clearly describe in Figure 275 that the Copy Command is referring to the NVM Command Set Copy command (not a Base command).

Remove incorrectly integrated, duplicate definition from Figure 280 NSFEAT field (bit 0 incorrectly duplicated the RMEDIA indication).

Update Feature value figure (Figure 316) to point to the appropriate command set as necessary for command set specific values.

The Keep Alive Timer value did not name the field that contains the timer granularity value (the KAS field). The name of the field was added to the description.

Use of the RTD3 Entry latency timer did not take CSTS.ST into account. Text was added.

NVM Express® NVM Command Set:

Added and corrected trademark attributions.

Correct definitions of User Data Read Access command and new Host Data Read command to align SMART and EG log page statistics for backward compatibility and match naming.

Clarify the distinction between the actual User Data format, and the Format Index used to specify which User Data format to use, rather than use the same term for both concepts.

Non-Volatility requirements were listed in an "informative" section; this was corrected.

The Deallocated or Unwritten Logical Block error was listed as a Command Specific Status Value for the Copy command. It is, however, a Media and Data Integrity Error Value, and the Copy command use of that error is already explained in Figure 16. The confusing text was removed.

Clarify that bits 3:0 of the FLBAS field are the LSB of the Format Index with which the namespace was formatted if there are > 16 formats supported.

Correct PIP field description that PI in the first bytes of the metadata was obsoleted in 2.0 (not in 1.4). Version 2.0 or later implementations that support PI are only allowed to support PI at the last bytes of the metadata.

In the description of NPWG, NPWA, NPDG, NPDA, and NOWS fields, the reference to the NSABP bit was a cut/paste error intended to fix the incorrect 1.4 reference to the OPTPEFT bit. To restore compatibility with 1.4, this was changed to reference the OPTPERF bit as originally intended.

NVM Express® Zoned Namespace Command Set:

Added and corrected trademark attributions.

Correct definitions of User Data Read Access command and new Host Data Read command to align SMART and EG log page statistics for backward compatibility and match naming.

Added definition of “write operation”, which had been accidentally deleted during integration.

Clarify the distinction between the actual User Data format, and the Format Index used to specify which User Data format to use, rather than use the same term for both concepts.

In the Sequential Write Required Zones section, replace usage of “User Data Out Command” with appropriate terms – which had been accidentally removed during integration.

In the ZSC to ZSIO state transition, clarify that a “successful operation” means that any data is written to the zone (and not necessarily that an actual write command completes successfully).

NVM Express® Key Value Command Set:

Added and corrected trademark attributions.

Correct definitions of User Data Read Access command and new Host Data Read command to align SMART and EG log page statistics for backward compatibility and match naming.

Clarify the distinction between the actual User Data format, and the Format Index used to specify which User Data format to use, rather than use the same term for both concepts.

NVM Express® PCI Express Transport:

Added and corrected trademark attributions.

NVM Express® RDMA Transport:

Added and corrected trademark attributions.

NVM Express® TCP Transport:

Added and corrected trademark attributions.

NVM Express® Management Interface:

Added and corrected trademark attributions. Added common Figure showing specification relationships.

Editor's Note:

BLACK text indicates unchanged text; **RED** text indicates new or deleted text; **GREEN** text indicates editor notes.

Description of NVM Express® Base specification changes

Modify a portion of the LEGAL DISCLAIMER as shown below:

LEGAL DISCLAIMER

...

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.

Modify a portion of section 1.1 as shown below:

1.1 Overview


The NVM Express® (NVMe®) interface allows host software to communicate with a non-volatile memory subsystem. This interface is optimized for all storage solutions, attached using a variety of transports including PCI Express®, Ethernet, InfiniBand™, and Fibre Channel. The mapping of extensions defined in this document to a specific NVMe Transport are defined in an NVMe Transport binding specification. The NVMe Transport binding specification for Fibre Channel is defined in INCITS 556 Fibre Channel – Non-Volatile Memory Express - 2 (FC-NVMe-2).

...

1.1.1 NVM Express® Specification Family

Figure 1 shows the relationship of the NVM Express specifications to each other within the NVMe™ family of specifications.

Figure 1: NVMe Family of Specifications



...

Modify a portion of section 1.6 as shown below:

1.6 I/O Command Set specific definitions used in the NVMe Base specification

The following terms used in this specification are defined in each I/O Command Set specification.

1.6.1 User Data Format

An I/O Command Set specific format ~~used to that~~ describes the layout of the data on the NVM media.

1.6.2 ~~User Data Read Access~~ SMART Data Units Read Command

An I/O Command Set specific command that results in the controller reading user data, but may or may not return the data to the host.

1.6.3 ~~User Data In Endurance Group Host Read~~ Command

An I/O Command Set specific command that ~~transfers user data from the controller to the host~~ results in the controller reading user data, but may or may not return the data to the host.

1.6.4 User Data Out Command

An I/O Command Set specific command that results in the controller writing user data, but may or may not transfer user data from the host to the controller.

1.6.5 SMART Host Read Command

An I/O Command Set specific command that results in the controller reading user data, but may or may not return the data to the host.

1.6.6 Format Index

A value used to index into the I/O Command Set Specific Format table (i.e., the User Data Format number).

Modify a portion of section 4.2 as shown below:

4.2 Feature Values

The Get Features command (refer to section 5.15) and Set Features command (refer to section 5.27) may be used to read and modify operating parameters of the controller. The operating parameters are grouped and identified by Feature Identifiers. Each Feature Identifier contains one or more attributes that may affect the behavior of the Feature.

If bit 4 is set to '1' in the Optional NVM Command Support (ONCS) field of the Identify Controller data structure in Figure 275, then for each Feature, there are three settings: default, saved, and current. If bit 4 is cleared to '0' in the Optional NVM Command Support field of the Identify Controller data structure, then the controller only supports a current and default value for each Feature. In this case, the current value may be persistent across power cycles and resets based on the information specified in Figure 316.

If bit 4 is set to '1' in the ONCS field, then ~~E~~each Feature has supported capabilities (refer to Figure 195), which are discovered ~~by~~ using the Supported Capabilities value in the Select field in Get Features (refer to Figure 192).

The default value for each Feature is vendor specific and set by the manufacturer unless otherwise specified. The default value is not changeable.

...

Modify a portion of section 5.14 as shown below:

5.14 Format NVM command

...

Figure 189: Format NVM – Command Dword 10

Bits	Description
31:14	Reserved
13:12	<p>LBA Format Upper (LBAFU): This field specifies the most significant 2 bits of the Format Index of the User Data Format to apply to the NVM media. This corresponds to the User Data Formats indicated in the Identify command, refer to the Identify Namespace data structure and the LBA Format data structure in the applicable I/O Command Set specification. Only supported User Data Formats shall be selected. If an unsupported User Data Format is selected, the controller shall abort the command with a status code of Invalid Format.</p> <p>This field is ignored if if the host has cleared the LBA Format Extension Enable (LBAFEE) field is cleared to 0h in the Host Behavior Support feature (refer to section 5.27.1.18).</p> <p>NOTE: This field applies to all User Data Formats. The original name has been retained for historical continuity.</p>
...	
03:00	<p>LBA Format Lower (LBAFL): This field specifies the least significant 4 bits of the Format Index of the User Data Format to apply to the NVM media. This corresponds to the User Data Formats indicated in the Identify Namespace data structure, refer to the Identify Namespace data structure and the I/O Command Set specific Format Data Structure in the applicable I/O Command Set specification. Only supported User Data Formats shall be selected. If an unsupported User Data Format is selected, the controller shall abort the command with a status code of Invalid Format.</p> <p>NOTE: This field applies to all User Data Formats. The original name has been retained for historical continuity.</p>
...	

...

Modify a portion of section 5.16 as shown below:

5.16 Get Log Page command

...

5.16.1 Log Specific Information

...

Figure 202: Get Log Page – Log Page Identifiers

Log Identifier	Scope	Log Page Name	Reference Section
...			
0Eh	Refer to the NVM I/O Command Set Specific		
...			
81h	NVM subsystem	Sanitize Status	5.16.1.25
82h to BFh	I/O Command Set Specific		
BFh	Refer to the Zoned Namespace Command Set		
C0h to FFh	Vendor specific ⁵		

...

5.16.1.3 SMART / Health Information (Log Identifier 02h)

...

If the log page is not supported on a per namespace basis, specifying a namespace identifier other than 0h or FFFFFFFFh should abort the command with a status code of Invalid Field in Command. If the controller does not abort the command, then the controller returns the controller log page. There is no namespace specific information defined in the SMART / Health Information log page in this revision of the specification, thus the controller log page and namespaces specific log page contain identical information.

...

Figure 207: Get Log Page – SMART / Health Information Log

Bytes	Description
...	
47:32	<p>Data Units Read: Contains the number of 512 byte data units the host has read from the controller as part of processing a User Data Read Access SMART Data Units Read Command; this value does not include metadata. This value is reported in thousands (i.e., a value of 1 corresponds to 1,000 units of 512 bytes read) and is rounded up (e.g., one indicates the that the number of 512 byte data units read is from 1 to 1,000, three indicates that the number of 512 byte data units read is from 2,001 to 3,000).</p> <p>Refer to the specific I/O Command Set specification for the list of User Data Read Access SMART Data Units Read Commands that affect this field.</p> <p>A value of 0h in this field indicates that the number of Data Units Read SMART Data Units Read is not reported.</p>
63:48	<p>Data Units Written: Contains the number of 512 byte data units the host has written to the controller as part of processing a User Data Out Command; this value does not include metadata. This value is reported in thousands (i.e., a value of 1 corresponds to 1,000 units of 512 bytes written) and is rounded up (e.g., one indicates that the number of 512 byte data units written is from 1 to 1,000, three indicates that the number of 512 byte data units written is from 2,001 to 3,000).</p> <p>Refer to the specific I/O Command Set specification for the list of User Data Out Commands that affect this field.</p> <p>A value of 0h in this field indicates that the number of Data Units Written is not reported.</p>
79:64	<p>Host Read Commands: Contains the number of User Data Read Access SMART Host Read Commands completed by the controller.</p> <p>Refer to the specific I/O Command Set specification for the list of User Data Read Access SMART Host Read Commands that affect this field.</p>
95:80	<p>Host Write Commands: Contains the number of User Data Out Commands completed by the controller.</p> <p>Refer to the specific I/O Command Set specification for the list of User Data Out Commands that affect this field.</p>
...	

...

5.16.1.8 Telemetry Host-Initiated (Log Identifier 07h)

This log consists of a header describing the log and zero or more Telemetry Data Blocks (refer to section 8.24). All Telemetry Data Blocks are 512 bytes in size. The controller shall initiate a capture of the controller's internal controller state to this log if the controller processes a Get Log Page command for this log with the Create Telemetry Host-Initiated Data bit set to '1' in the Log Specific field. If the host specifies a Log Page Offset Lower value that is not a multiple of 512 bytes in the Get Log Page command for this log, then the controller shall return an error with a status code set to ~~ef~~ Invalid Field in Command. This log page is global to the controller or global to the NVM subsystem.

...

5.16.1.10 Endurance Group Information (Log Identifier 09h)

...

Figure 217: Get Log Page – Endurance Group Information Log (Log Identifier 09h)

Bytes	Description
...	
111:96	Host Read Commands: Contains the number of User Data In Endurance Group Host Read Commands completed by the controller. Refer to the specific I/O Command Set specification for the list of User Data In Endurance Group Host Read Commands that affect this field.
127:112	Host Write Commands: Contains the number of User Data Out Commands completed by the controller. Refer to the specific I/O Command Set specification for the list of User Data Out Commands that affect this field.
...	

...

5.16.1.13 Asymmetric Namespace Access (Log Identifier 0Ch)

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

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

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

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

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

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

...

Modify a portion of section 5.17 as shown below:

5.17 Identify command

5.17.1 Identify command overview

...

Figure 273: Identify – CNS Values

CNS Value	O/M ¹	Definition	NSID ²	CNTID ³	CSI ⁴	Reference Section
Active Namespace Management						
...						
08h	M	I/O Command Set Independent Identify Namespace data structure.	Y	N	N	5.17.2.8
09h to 0Fh		Reserved				
...						
10h	O ⁵	Allocated Namespace ID list.	Y	N	N	5.17.2.8 5.17.2.9
11h	O ^{5, 11}	Identify Namespace data structure for the specified allocated NSID.	Y	N	N ⁸	NVM Command Set Specification 5.17.2.10
...						
16h	O ¹¹	A Namespace Granularity List (refer to the NVM Command Set Specification) is returned to the host for up to sixteen Namespace Granularity Entries.	N	N	N ⁸	NVM Command Set Specification 5.17.2.15
...						

...

5.17.2 Identify Data Structures

5.17.2.1 Identify Controller data structure (CNS 01h)

The Identify Controller data structure (refer to Figure 275) is returned to the host for the controller processing the command.

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

Bytes	I/O ¹	Admin ¹	Disc ¹	Description
Controller Capabilities and Features				
...				
295:280	O	O	R	Total NVM Capacity (TNVMCAP): This field indicates the total NVM capacity that is accessible by the controller. The value is in bytes. This field shall be supported if the Namespace Management capability (refer to section 8.11) is supported or if the Capacity Management capability (refer to section 8.3) is supported. Refer to section 3.8.
311:296	O	O	R	Unallocated NVM Capacity (UNVMCAP): This field indicates the unallocated NVM capacity that is accessible by the controller. The value is in bytes. This field shall be supported if the Namespace Management capability (refer to section 8.11) is supported or if the Capacity Management capability (refer to section 8.3) is supported. Refer to section 3.8.
...				
521:520	M	M	R	Optional NVM Command Support (ONCS): This field indicates the optional I/O commands and features supported by the controller. Refer to section 3.1.2. Bits 15:9 are reserved. Bit 8 if set to '1', then the controller supports the NVM Command Set Copy command. If cleared to '0', then the controller does not support the NVM Command Set Copy command. ...
...				

...

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

...

Figure 280: Identify – I/O Command Set Independent Identify Namespace Data Structure

Bytes	O/M ¹	Description
00	M	<p>Common Namespace Features (NSFEAT): This field defines features of the namespace.</p> <p>Bits 7:5 are reserved.</p> <p>Bit 4 Rotational Media (RMEDIA) if set to '1' indicates that the namespace stores data on rotational media (refer to section 8.20). If cleared to '0', indicates that the namespace does not store data on rotational media.</p> <p>Bit 3 (UIDREUSE) if set to '1' indicates that the value in the NGUID field for this namespace, if non-zero, is never reused by the controller and that the value in the EUI64 field for this namespace, if non-zero, is never reused by the controller. If cleared to '0', then the NGUID value may be reused and the EUI64 value may be reused by the controller for a new namespace created after this namespace is deleted. This bit shall be cleared to '0' if both NGUID and EUI64 fields are cleared to 0h. Refer to section 4.5.1.</p> <p>Bit 2:40 are reserved.</p> <p>Bit 0 (Rotational Media) if set to '1' indicates that the namespace is associated with rotational media (refer to section 8.20). If cleared to '0', then the namespace is not associated with rotational media.</p>
...		
<p>NOTES:</p> <p>1. O/M definition: O = Optional, M = Mandatory.</p>		

...

Modify a portion of section 5.27 as shown below:

5.27 Set Features command

...

5.27.1 Feature Specific Information

...

Figure 316: Set Features – Feature Identifiers

Feature Identifier	Current Setting Persists Across Power Cycle and Reset ²	Uses Memory Buffer for Attributes	Feature Name
00h	Reserved		
01h	No	No	Arbitration
02h	No	No	Power Management
03h	Refer to the NVM Command Set		
04h	No	No	Temperature Threshold
05h	Refer to the NVM Command Set		
06h	No	No	Volatile Write Cache
...			

Figure 316: Set Features – Feature Identifiers

Feature Identifier	Current Setting Persists Across Power Cycle and Reset ²	Uses Memory Buffer for Attributes	Feature Name
09h	No	No	Interrupt Vector Configuration
0Ah	Refer to the NVM Command Set		
0Bh	No	No	Asynchronous Event Configuration
...			
14h	No	No	Predictable Latency Mode Window
15h	Refer to the NVM Command Set		
16h	No	Yes	Host Behavior Support
...			
20h	Refer to the Key Value Command Set-Specific		
...			

...

5.27.1.12 Keep Alive Timer (Feature Identifier 0Fh)

...

This Feature controls the Keep Alive Timer. Refer to section 3.9 for Keep Alive details. The attributes are specified in Command Dword 11.

If a Get Features command is submitted for this Feature, the attributes specified in Figure 341 are returned in Dword 0 of the completion queue entry for that command.

Figure 341: Keep Alive Timer – Command Dword 11

Bits	Description
31:00	<p>Keep Alive Timeout (KATO): This field specifies the timeout value for the Keep Alive feature in milliseconds. The controller rounds up the value specified to the granularity indicated in the KAS field in the Identify Controller data structure. If cleared to 0h, then the Keep Alive Timer is disabled.</p> <p>The default value for this field is 0h for PCIe and other NVMe transports that do not require use of the Keep Alive feature (e.g., NVMe over PCIe). For NVMe transports that require use of the Keep Alive feature (e.g., RDMA and TCP), the default value for this field is 1D4C0h (i.e., 120,000 milliseconds or 2 minutes) rounded up to that the granularity indicated in the KAS field.</p> <p>Refer to the applicable NVMe Transport Binding specification for details.</p>

...

Modify a portion of section 8.15 as shown below:

8.15 Power Management

...

8.15.4 Runtime D3 Transitions

...

The latency reported is based on a normal shutdown with optimal controller settings preceding the RTD3 resume. The latency reported assumes that host software enables and initializes the controller and sends a 4 KiB read operation.

If CSTS.ST is cleared to '0', then the RTD3 Entry Latency is the expected elapsed time from the time CC.SHN is set to 01b by host software until CSTS.SHST is set to 10b by the controller. When CSTS.SHST is set to 10b, it is safe for host software to remove power from the controller.

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

Description of NVM Express® NVM Command Set specification changes

Modify a portion of the LEGAL DISCLAIMER as shown below:

LEGAL DISCLAIMER

...

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.
PCIe® is a registered trademark of PCI-SIG.

Modify a portion of section 1.2 as shown below:

1.2 Scope

Figure 1 shows the relationship of the NVM Express® NVM Command Set Specification to other specifications within the NVMe Family of Specifications.

Figure 1: NVMe Family of Specifications



Modify a portion of section 1.4 as shown below:

1.4 Definitions

...

1.4.2.1 ~~User Data Read Access~~ SMART Data Units Read Command

The Compare command, ~~Copy command~~, Read command, and Verify command.

1.4.2.2 User Data Format

~~The LBA Format is the data format of the NVM Command Set. The layout of the data on the NVM media as described by the LBA Format of the namespace.~~

1.4.2.3 ~~User Data In~~ Endurance Group Host Read Command

The Compare command, Copy command, ~~and~~ Read command, ~~and~~ Verify command.

1.4.2.4 User Data Out Command

The Copy command and Write command.

1.4.2.5 SMART Host Read Command

The Compare command, Copy command, and Read command.

1.4.2.6 Format Index

A value used to index into the LBA Format data structure and the Extended LBA Format data structure.

...

Modify a portion of section 2.1 as shown below:

2.1 Theory of operation

...

2.1.4.2.1 AWUPF/NAWUPF Example (Informative)

...

If the size of write command A is larger than the AWUPF/NAWUPF value, then there is no guarantee of the state of the data contained in the specified LBA range after the power fail or error condition.

2.1.4.2.2 Non-volatile requirements

After a write command has completed, reads for that location which are subsequently submitted shall return the data from that write command and not an older version of the data from previous write commands with the following exception:

...

Modify a portion of section 3.2 as shown below:

3.2 NVM Command Set Commands

...

Figure 18: Opcodes for NVM Commands

Opcode by Field			Combined Opcode ¹	Command ²	Reference Section
(07)	(06:02)	(01:00)			
Standard Command	Function	Data Transfer ³			
...					
0b	001 01b	01b	15h	Reservation Release	Base
0b	001 10b	01b	19h	Copy	3.2.2 Base
Vendor Specific					
...					

...

3.2.2 Copy command

...

3.2.2.1 Command Completion

...

Figure 37: Copy – Command Specific Status Values

Value	Description
...	
87h	Deallocated or Unwritten Logical Block: The command failed due to an attempt to copy from an LBA range containing a deallocated or unwritten logical block.

Modify a portion of section 4.1 as shown below:

4.1 Admin Command behavior for the NVM Command Set

...

4.1.2 Format NVM command

...

If ~~a host does not set~~ the LBA Format Extension Enable (LBAFEE) field ~~is not set~~ to 1h in the Host Behavior Support feature (refer to the Host Behavior Support section in the NVMe Base Specification), then ~~a the~~ controller aborts a Format NVM command with a status code of Invalid Namespace or Format that specifies a format ~~of~~ (refer to section 5.2) ~~of~~:

- a) 16b Guard Protection Information with the STS field set to a non-zero value;

- b) 32b Guard Protection Information; ~~and~~ or
- c) 64b Guard Protection Information.

Figure 78: Format NVM – Command Dword 10 – NVM Command Set Specific Fields

Bits	Description
08	Protection Information Location (PIL): If set to '1' and protection information is enabled (refer to section 5.1), then protection information is transferred as the first bytes of metadata. If cleared to '0' and protection information is enabled, then protection information is transferred as the last bytes of metadata. This setting is reported in the End-to-end Data Protection Type Settings (DPS) field of the Identify Namespace data structure and is constrained by the End-to-end Data Protection Capabilities (DPC) field of the Identify Namespace data structure. For implementations compliant with version 4-4 1.0 or later of the NVMe-Base NVMe Command Set Specification or later , this field shall be cleared to '0'.
...	

...

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

...

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

Bytes	O/M ¹	Description
...		
25	M	<p>Number of LBA Formats (NLBAF): This field defines the number of supported LBA data size and metadata size combinations supported by the namespace. LBA formats shall be allocated in order (starting with 0) and packed sequentially. This is a 0's based value. The maximum number of LBA formats that may be indicated as supported is:</p> <ul style="list-style-type: none"> a) 16 if the host has cleared the LBA Format Extension Enable (LBAFEE) field is cleared to 0h in the Host Behavior Support feature (refer to the Host Behavior Support section in the NVMe Base Specification); or b) 64 if the host has set the LBAFEE field is set to 1h in the Host Behavior Support feature (refer to the Host Behavior Support section in the NVMe Base Specification). <p>The supported LBA formats are indicated in bytes 128 to 383 in this data structure. The LBA Format fields with an index beyond the value set in this field are invalid and not supported. LBA Formats that are valid, but not currently available may be indicated by setting the LBA Data Size for that LBA Format to 0h.</p> <p>The metadata may be either transferred as part of the LBA (creating an extended LBA which is a larger LBA size that is exposed to the application) or may be transferred as a separate contiguous buffer of data. The metadata shall not be split between the LBA and a separate metadata buffer.</p> <p>It is recommended that software and controllers transition to an LBA size that is 4 KiB or larger for ECC efficiency at the controller. If providing metadata, it is recommended that at least 8 bytes are provided per logical block to enable use with end-to-end data protection, refer to section 5.8.3.</p>

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

26	M	<p>Formatted LBA Size (FLBAS): This field indicates the LBA data size & metadata size combination that the namespace has been formatted with (refer to section 4.1.2).</p> <p>Bits 7 is reserved.</p> <p>Bits 6:5 indicate the most significant 2 bits of the Format Index of the supported LBA Formats indicated in this data structure that was used to format the namespace. These bits are ignored if the number of supported LBA Formats is less than or equal to 16. If the NLBAF field is less than or equal to 16, then the host should ignore these bits.</p> <p>Bit 4 if set to '1' indicates that the metadata is transferred at the end of the data LBA, creating an extended data LBA. Bit 4 if cleared to '0' indicates that all of the metadata for a command is transferred as a separate contiguous buffer of data. Bit 4 is not applicable when there is no metadata.</p> <p>Bits 3:0 indicates the least significant 4 bits of the Format Index one of the 16 supported LBA Formats indicated in this data structure that was used to format the namespace.</p>																				
...																						
29	M	<p>End-to-end Data Protection Type Settings (DPS): This field indicates the protection information Type settings for the end-to-end data protection feature. Refer to section 5.1.</p> <table border="1" data-bbox="532 695 1458 1178"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>7:4</td> <td>Reserved</td> </tr> <tr> <td>3</td> <td> <p>Protection Information Position (PIP): This bit indicates that the protection information, if enabled, is transferred as the first bytes of metadata. Bit 3 if cleared to '0' indicates that the protection information, if enabled, is transferred as the last bytes of metadata. For implementations compliant to version 1.4 1.0 or later of the NVMe Base NVM Command Set Specification or later, this bit shall be cleared to '0'.</p> </td> </tr> <tr> <td>2:0</td> <td> <p>Protection Information Type (PIT): This field indicates whether protection information is enabled and the type of protection information enabled. The values for this field have the following meanings:</p> <table border="1" data-bbox="667 1003 1425 1178"> <thead> <tr> <th>Value</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>Protection information is not enabled</td> </tr> <tr> <td>001b</td> <td>Protection information is enabled, Type 1</td> </tr> <tr> <td>010b</td> <td>Protection information is enabled, Type 2</td> </tr> <tr> <td>011b</td> <td>Protection information is enabled, Type 3</td> </tr> <tr> <td>100b to 111b</td> <td>Reserved</td> </tr> </tbody> </table> </td> </tr> </tbody> </table>	Bits	Description	7:4	Reserved	3	<p>Protection Information Position (PIP): This bit indicates that the protection information, if enabled, is transferred as the first bytes of metadata. Bit 3 if cleared to '0' indicates that the protection information, if enabled, is transferred as the last bytes of metadata. For implementations compliant to version 1.4 1.0 or later of the NVMe Base NVM Command Set Specification or later, this bit shall be cleared to '0'.</p>	2:0	<p>Protection Information Type (PIT): This field indicates whether protection information is enabled and the type of protection information enabled. The values for this field have the following meanings:</p> <table border="1" data-bbox="667 1003 1425 1178"> <thead> <tr> <th>Value</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>Protection information is not enabled</td> </tr> <tr> <td>001b</td> <td>Protection information is enabled, Type 1</td> </tr> <tr> <td>010b</td> <td>Protection information is enabled, Type 2</td> </tr> <tr> <td>011b</td> <td>Protection information is enabled, Type 3</td> </tr> <tr> <td>100b to 111b</td> <td>Reserved</td> </tr> </tbody> </table>	Value	Definition	000b	Protection information is not enabled	001b	Protection information is enabled, Type 1	010b	Protection information is enabled, Type 2	011b	Protection information is enabled, Type 3	100b to 111b	Reserved
Bits	Description																					
7:4	Reserved																					
3	<p>Protection Information Position (PIP): This bit indicates that the protection information, if enabled, is transferred as the first bytes of metadata. Bit 3 if cleared to '0' indicates that the protection information, if enabled, is transferred as the last bytes of metadata. For implementations compliant to version 1.4 1.0 or later of the NVMe Base NVM Command Set Specification or later, this bit shall be cleared to '0'.</p>																					
2:0	<p>Protection Information Type (PIT): This field indicates whether protection information is enabled and the type of protection information enabled. The values for this field have the following meanings:</p> <table border="1" data-bbox="667 1003 1425 1178"> <thead> <tr> <th>Value</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>Protection information is not enabled</td> </tr> <tr> <td>001b</td> <td>Protection information is enabled, Type 1</td> </tr> <tr> <td>010b</td> <td>Protection information is enabled, Type 2</td> </tr> <tr> <td>011b</td> <td>Protection information is enabled, Type 3</td> </tr> <tr> <td>100b to 111b</td> <td>Reserved</td> </tr> </tbody> </table>	Value	Definition	000b	Protection information is not enabled	001b	Protection information is enabled, Type 1	010b	Protection information is enabled, Type 2	011b	Protection information is enabled, Type 3	100b to 111b	Reserved									
Value	Definition																					
000b	Protection information is not enabled																					
001b	Protection information is enabled, Type 1																					
010b	Protection information is enabled, Type 2																					
011b	Protection information is enabled, Type 3																					
100b to 111b	Reserved																					
...																						
65:64	O	<p>Namespace Preferred Write Granularity (NPWG): This field indicates the smallest recommended write granularity in logical blocks for this namespace. This is a 0's based value. If the NSABP OPTPERF bit is cleared to '0', then this field is reserved.</p> <p>The size indicated should be less than or equal to Maximum Data Transfer Size (MDTS) that is specified in units of minimum memory page size. The value of this field may change if the namespace is reformatted. The size should be a multiple of Namespace Preferred Write Alignment (NPWA).</p> <p>Refer to section 5.8.2 for how this field is utilized to improve performance and endurance.</p>																				
67:66	O	<p>Namespace Preferred Write Alignment (NPWA): This field indicates the recommended write alignment in logical blocks for this namespace. This is a 0's based value. If the NSABP OPTPERF bit is cleared to '0', then this field is reserved.</p> <p>The value of this field may change if the namespace is reformatted.</p> <p>Refer to section 5.8.2 for how this field is utilized to improve performance and endurance.</p>																				
69:68	O	<p>Namespace Preferred Deallocate Granularity (NPDG): This field indicates the recommended granularity in logical blocks for the Dataset Management command with the Attribute – Deallocate bit set to '1' in Dword 11. This is a 0's based value. If the NSABP OPTPERF bit is cleared to '0', then this field is reserved.</p> <p>The value of this field may change if the namespace is reformatted. The size should be a multiple of Namespace Preferred Deallocate Alignment (NPDA).</p> <p>Refer to section 5.8.2 for how this field is utilized to improve performance and endurance.</p>																				

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

71:70	O	<p>Namespace Preferred Deallocate Alignment (NPDA): This field indicates the recommended alignment in logical blocks for the Dataset Management command with the Attribute – Deallocate bit set to ‘1’ in Dword 11. This is a 0’s based value. If the NSABP OPTPERF bit is cleared to ‘0’, then this field is reserved.</p> <p>The value of this field may change if the namespace is reformatted.</p> <p>Refer to section 5.8.2 for how this field is utilized to improve performance and endurance.</p>
73:72	O	<p>Namespace Optimal Write Size (NOWS): This field indicates the size in logical blocks for optimal write performance for this namespace. This is a 0’s based value. If the NSABP OPTPERF bit is cleared to ‘0’, then this field is reserved.</p> <p>The size indicated should be less than or equal to Maximum Data Transfer Size (MDTS) that is specified in units of minimum memory page size. The value of this field may change if the namespace is reformatted. The value of this field should be a multiple of Namespace Preferred Write Granularity (NPWG).</p> <p>If the namespace is associated with an NVM set, NOWS defined for this namespace shall be set to the Optimal Write Size field setting defined in NVM Set Attributes Entry (refer to the Namespace Identification Descriptor in the NVMe Base Specification) for the NVM Set with which this namespace is associated. If NOWS is not supported, the Optimal Write Size field in NVM Sets Attributes Entry (refer to the Namespace Identification Descriptor in the NVMe Base Specification) for the NVM Set with which this namespace is associated should be used by the host for I/O optimization.</p> <p>Refer to section 5.8.2 for how this field is utilized to improve performance and endurance.</p>
...		
LBA Formats		
131:128	M	<p>LBA Format 0 Support (LBAF0): This field indicates the LBA format 0 that is supported by the controller. The LBA format field is defined in Figure 97.</p> <p>Additional information may be provided in the ELBAF0 field (refer to Figure 99).</p>
...		

...

4.1.5.3 I/O Command Set Specific Identify Namespace Data Structure (CNS 05h)

...

Figure 100: NVM Command Set I/O Command Set Specific Identify Namespace Data Structure (CSI 00h)

Bytes	O/M ¹	Description
7:0	O	<p>Logical Block Storage Tag Mask (LBSTM): Identifies the mask for the Storage Tag field for the protection information (refer to section 5.1). The size of the mask contained in this field is defined by the STS field. If the size of the mask contained in this field is less than 64 bits, the mask is contained in the least-significant bits of this field.</p> <p>If end-to-end protection is not enabled in the namespace, then this field is ignored.</p> <p>If:</p> <ul style="list-style-type: none"> a) end-to-end protection is enabled; b) 16b Guard Protection Information format is used; and c) the 16BPISTM bit is set to ‘1’, <p>then all bits in the mask shall be set to ‘1’.</p>

Figure 100: NVM Command Set I/O Command Set Specific Identify Namespace Data Structure (CSI 00h)

Bytes	O/M ¹	Description								
8	O	<p>Protection Information Capabilities (PIC): This field indicates the capabilities for the protection information formats.</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>7:2</td> <td>Reserved</td> </tr> <tr> <td>1</td> <td> <p>16b Guard Protection Information Storage Tag Mask (16BPISM): If set to '1', then the LBSTM field shall have all bits set to '1' for the 16b Guard Protection Information. If cleared to '0', then the Logical Block Storage Tag Mask field is allowed to have any bits set to '1' for the 16b Guard Protection Information.</p> </td> </tr> <tr> <td>0</td> <td> <p>16b Guard Protection Information Storage Tag Support (16BPISTS): If set to '1', then the end-to-end protection 16b Guard Protection Information format (refer to section 5.2.1.1) supports a non-zero value in the STS field. If cleared to '0', then the end-to-end protection 16b Guard Protection Information format support requires that the STS field be cleared to 0h (i.e., the Storage Tag field is not supported).</p> <p>If the 32b Guard Protection Information or 64b Guard Protection Information is supported in any LBA format (refer to Figure 97 and Figure 100), then this bit shall be set to '1'.</p> </td> </tr> </tbody> </table>	Bits	Description	7:2	Reserved	1	<p>16b Guard Protection Information Storage Tag Mask (16BPISM): If set to '1', then the LBSTM field shall have all bits set to '1' for the 16b Guard Protection Information. If cleared to '0', then the Logical Block Storage Tag Mask field is allowed to have any bits set to '1' for the 16b Guard Protection Information.</p>	0	<p>16b Guard Protection Information Storage Tag Support (16BPISTS): If set to '1', then the end-to-end protection 16b Guard Protection Information format (refer to section 5.2.1.1) supports a non-zero value in the STS field. If cleared to '0', then the end-to-end protection 16b Guard Protection Information format support requires that the STS field be cleared to 0h (i.e., the Storage Tag field is not supported).</p> <p>If the 32b Guard Protection Information or 64b Guard Protection Information is supported in any LBA format (refer to Figure 97 and Figure 100), then this bit shall be set to '1'.</p>
		Bits	Description							
		7:2	Reserved							
		1	<p>16b Guard Protection Information Storage Tag Mask (16BPISM): If set to '1', then the LBSTM field shall have all bits set to '1' for the 16b Guard Protection Information. If cleared to '0', then the Logical Block Storage Tag Mask field is allowed to have any bits set to '1' for the 16b Guard Protection Information.</p>							
0	<p>16b Guard Protection Information Storage Tag Support (16BPISTS): If set to '1', then the end-to-end protection 16b Guard Protection Information format (refer to section 5.2.1.1) supports a non-zero value in the STS field. If cleared to '0', then the end-to-end protection 16b Guard Protection Information format support requires that the STS field be cleared to 0h (i.e., the Storage Tag field is not supported).</p> <p>If the 32b Guard Protection Information or 64b Guard Protection Information is supported in any LBA format (refer to Figure 97 and Figure 100), then this bit shall be set to '1'.</p>									
11:9	Reserved									
Extended LBA Format										
15:12	O	<p>Extended LBA Format 0 Support (ELBAF0): This field indicates additional LBA Format 0 information related to the LBA Format 0 Support (LBAF0) field in the Identify Namespace data structure. The Extended LBA format field is defined in Figure 101.</p>								
19:16	O	<p>Extended LBA Format 1 Support (ELBAF1): This field indicates additional LBA Format 1 information related to the LBA Format 1 Support (LBAF1) field in the Identify Namespace data structure. The Extended LBA format field is defined in Figure 101.</p>								
...								
267:264	O	<p>Extended LBA Format 63 Support (ELBAF63): This field indicates additional LBA Format 63 information related to the LBA Format 63 Support (LBAF63) field in the Identify Namespace data structure. The Extended LBA format field is defined in Figure 101.</p>								
4095:268	O	Reserved								
<p>NOTES:</p> <p>1. O/M definition: O = Optional, M = Mandatory.</p>										

4.1.5.6 Namespace Granularity List (CNS 16h)

If the controller supports reporting of Namespace Granularity (refer to section 5.3), then a Namespace Granularity List (refer to Figure 102) is returned to the host for up to:

- a) 16 namespace granularity descriptors (refer to Figure 103) if the ~~host has cleared the~~ LBA Format Extension Enable (LBAFEE) field ~~is cleared~~ to 0h in the Host Behavior Support feature (refer to the Host Behavior Support section in the NVMe Base Specification); or
- b) 64 namespace granularity descriptors if the ~~host has set the~~ LBAFEE field ~~is set~~ to 1h in the Host Behavior Support feature.

...

4.1.6 Namespace Management command

...

If ~~a host does not set~~ the LBA Format Extension Enable (LBAFEE) field ~~is not set~~ to 1h in the Host Behavior Support feature (refer to the Host Behavior Support section in the NVMe Base Specification), then a controller aborts a Namespace Management command with a status code of Invalid Namespace or Format that specifies to create a namespace that is formatted with (refer to section 5.2):

...

Description of NVM Express® Zoned Namespace Command Set specification changes

Modify a portion of the LEGAL DISCLAIMER as shown below:

LEGAL DISCLAIMER

...

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.
PCIe® is a registered trademark of PCI-SIG.

Modify a portion of section 1.2 as shown below:

1.2 Scope

Figure 1 shows the relationship of the NVM Express® Zoned Namespace Command Set Specification to other specifications within the NVMe Family of Specifications.

Figure 1: NVMe Family of Specifications



Modify a portion of section 1.4 as shown below:

1.4 Definitions

...

1.4.2.1 ~~User Data Read Access~~ SMART Data Units Read Command

A ~~User Data Read Access~~ SMART Data Units Read Command as defined in the NVM Command Set Specification.

~~1.4.2.2 User Data Format~~

~~The LBA Format (defined in the NVM Command Set specification) and the LBA Format Extension (defined in this specification).~~

1.4.2.3 ~~User Data In~~ Endurance Group Host Read Command

An ~~User Data In~~ Endurance Group Host Read Command as defined in the NVM Command Set Specification.

...

1.4.2.5 SMART Host Read Command

A SMART Host Read Command as defined in the NVM Command Set Specification.

1.4.2.6 Format Index

A value used to index into the LBA Format data structure, the Extended LBA Format data structure, and the LBA Format Extension data structure.

1.4.3 Definitions from the NVM Command Set specification

The following terms are defined in the NVM Command Set specification and are used in this specification:

- a) LBA range
- b) logical block
- c) logical block address (LBA)
- d) **User Data Format**

...

1.4.4.4 write operation

An operation initiated by a Write command, a Write Zeroes command, a Write Uncorrectable command, a Copy command, or a Zone Append command.

Modify a portion of section 1.5 as shown below:

1.5 References

NVMe[®] Base Specification, Revision 2.0. **dd-mm-2021** Available from <https://www.nvmexpress.org>.

NVMe[®] NVM Command Set Specification, Revision 1.0. **dd-mm-2021** Available from <https://www.nvmexpress.org>.

Modify a portion of section 2.1 as shown below:

2.1 Theory of operation

...

2.1.1.2.1.1 Writing in Sequential Write Required Zones

~~User Data Out commands, the Write Uncorrectable command, and the Write Zeroes command may be used to write user data to specific zones of zone type Sequential Write Required.~~

The following commands may be used to write to logical blocks in a specific zone of zone type Sequential Write Required:

- Write command;
- Write Zeroes command;
- Write Uncorrectable command;
- Copy command; and
- Zone Append command.

A write pointer is maintained for each zone in the zoned namespace that indicates the next writable logical block address in that zone. The write pointer is valid for a subset of the zone states as defined in Figure 4.

...

The write pointer for a zone in the ZSE:Empty state, the ZSIO:Implicitly Opened state, the ZSEO:Explicitly Opened state, or the ZSC:Closed state shall be increased by the number of logical blocks written on successful completion of a write operation.

If the controller is not able to successfully write to all logical blocks specified by a ~~User Data Out Command~~ that initiates a write operation, then the write pointer shall:

- a) be set to a value within the range of LBAs specified in that ~~User Data Out Command~~;
- b) be set to one greater than the last LBA in the range of LBAs specified in that ~~User Data Out Command~~; or
- c) become invalid (i.e., due to transitioning to the ZSRO:Read Only state or the ZSO:Offline state, or due to a Zone Active Excursion (refer to section 5.6)).

...

The controller shall abort a ~~User Data Out Command~~ that writes to a zone that is in the ZSF:Full state with a status code of Zone Is Full.

The controller shall abort a ~~User Data Out Command~~ that writes to a zone that is in the ZSRO:Read Only state, with a status code of Zone Is Read Only.

The controller shall abort a ~~User Data Out Command~~ that writes to a zone that is in the ZSO:Offline state, with a status code of Zone Is Offline.

...

The controller shall abort a ~~User Data Out Command~~ that initiates a write operation that has a starting LBA in one zone, and for which the Number of Logical blocks exceeds the remaining number of logical blocks in that zone, with a status code of Zone Boundary Error.

...

2.1.1.3.4 ZSC:Closed state

...

Transition ZSC:ZSIO: The zone shall transition from the ZSC:Closed state to the ZSIO:Implicitly Opened state, if there are available resources as defined in section 2.1.1.4 and a write operation that writes one or more logical blocks to the of that zone completes successfully.

...

Modify a portion of section 4.1 as shown below:

4.1 Admin Command behavior for the Zoned Namespace Command Set

...

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

...

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

Bytes	O/M ¹	Description
...		
LBA Format Extensions		

Bytes	O/M ¹	Description
2831:2816	M	LBA Format 0 Extension (LBAFE0): This field indicates the LBA format Extension 0 that is supported by the controller. The Zone format field is defined in Figure 47.
...		

Modify a portion of sections A.3 and A.4 as shown below:

A.3 Open Zone Considerations

A.3.1. Overview

...

A.3.2. Zones in the ZSEO:Explicitly Opened Zones and the ZSIO:Implicitly Opened states

...

A.3.3. Opening and Closing Zones

...

A.3.4. Zone Send Action of Finish Zone Considerations

...

A.4 Partial Failures

...

A.4.1. Overview

...

Description of NVM Express® Key Value Command Set specification changes

Modify a portion of the LEGAL DISCLAIMER as shown below:

LEGAL DISCLAIMER

...

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.
PCIe® is a registered trademark of PCI-SIG.

Modify a portion of section 1.2 as shown below:

1.2 Scope

Figure 1 shows the relationship of the NVM Express® Key Value Command Set Specification to other specifications within the NVMe Family of Specifications.

Figure 1: NVMe Family of Specifications



Modify a portion of section 1.4 as shown below:

1.4 Definitions

...

1.4.2.1 User Data Format

~~The Key Value Format is the data format of the Key Value Command Set. The layout of the data on the NVM media as described by the Key Value Format of the namespace.~~

1.4.2.2 ~~User Data In Endurance Group Host Read Command~~

A Retrieve command.

...

1.4.2.4 ~~User Data Read Access SMART Data Units Read Command~~

A Retrieve command.

1.4.2.5 SMART Host Read Command

A Retrieve command.

1.4.2.6 Format Index

A value used to index into the KV Format data structure.

Modify a portion of section 4.1 as shown below:

4.1 Admin Command behavior for the Key Value Command Set

...

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

...

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

Bytes	O/M ¹	Description
...		
KV Formats		
87:72	M	KV Format 0 Support (KVF0): This field indicates the KV format 0 that is supported by the controller. The KV format field is defined in Figure 37.
...		

...

Description of NVM Express® NVMe® over PCIe® Transport specification changes

Modify a portion of the LEGAL DISCLAIMER as shown below:

LEGAL DISCLAIMER

...

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.

Modify a portion of section 1.2 as shown below:

1.2 Scope

Figure 1 shows the relationship of the ~~NVM Command Set~~ NVM Express® NVMe® over PCIe® Transport Specification to other specifications within the NVMe Family of Specifications.

Figure 1: NVMe Family of Specifications



Modify a portion of section 3.8 as shown below:

3.8 Transport Specific Content

...

3.8.1.2 Offset 04h: CMD - Command

Figure 12: Offset 04h: CMD - Command

Bits	Type	Reset	Description
15:11	RO	0h	Reserved by PCI-SIG®
...			

Description of NVM Express® RDMA Transport specification changes

Modify a portion of the LEGAL DISCLAIMER as shown below:

LEGAL DISCLAIMER

...

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.
PCIe® is a registered trademark of PCI-SIG.

Modify a portion of section 1.2 as shown below:

1.2 Scope

Figure 1 shows the relationship of the ~~NVM Command Set~~ NVM Express® RDMA Transport Specification to other specifications within the NVMe Family of Specifications.

Figure 1: NVMe Family of Specifications



Description of NVM Express® TCP Transport specification changes

Modify a portion of the LEGAL DISCLAIMER as shown below:

LEGAL DISCLAIMER

...

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.
PCIe® is a registered trademark of PCI-SIG.

Modify a portion of section 1.2 as shown below:

1.2 Scope

Figure 1 shows the relationship of the ~~NVM Command Set~~ NVM Express® TCP Transport Specification to other specifications within the NVMe Family of Specifications.

Figure 1: NVMe Family of Specifications



Description of NVM Express® Management Interface specification changes

Modify a portion of the LEGAL DISCLAIMER as shown below:

LEGAL DISCLAIMER

...

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.

Modify a portion of section 1.1 as shown below:

1.1 Overview

The NVM Express® Management Interface Specification is a member of the NVMe Family of Specifications displayed in Figure 0TBD.

Figure 0TBD: NVMe Family of Specifications



The NVM Express® (NVMe®) interface ~~is a register-level interface that~~ allows in-band host software to communicate with an NVM Subsystem. Since this specification builds on the NVM Express® Base Specification, knowledge of the NVM Express® Base Specification is assumed.

...

Modify a portion of section 1.2 as shown below:

1.2 Scope

...

1.2.1 Outside of Scope

...

This specification does not cover management of non-transparent bridges or PCIe[®] switches. Coordination between multiple Requesters or a Requester and a device other than a Responder is outside the scope of this specification. Refer to section 1.8 for the definitions of Requester and Responder.

...

An enclosure may support comprehensive management capabilities using SCSI Enclosure Services, basic management capabilities using transport specific mechanisms, or no management capabilities. An example of basic enclosure management capabilities is Native PCIe Enclosure Management (NPEM) specified by the PCI-SIG[®] for PCI Express[®]. The specification of such transport specific basic management capabilities is outside the scope of this specification. This specification only defines comprehensive management using SCSI Enclosure Services.

...