



LEGAL NOTICE:

© Copyright 2007 to 2020 NVM Express™, Inc. ALL RIGHTS RESERVED.

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

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

NOTICE TO NON-MEMBERS OF NVM EXPRESS, INC.: If you are not a Member of NVM Express, Inc. and you have obtained a copy of this document, you only have a right to review this document or make reference to or cite this document. Any such references or citations to this document must acknowledge NVM Express, Inc. copyright ownership of this document. The proper copyright citation or reference is as follows: "© 2007 to 2020 NVM Express, Inc. ALL RIGHTS RESERVED." When making any such citations or references to this document you are not permitted to revise, alter, modify, make any derivatives of, or otherwise amend the referenced portion of this document in any way without the prior express written permission of NVM Express, Inc. Nothing contained in this document shall be deemed as granting you any kind of license to implement or use this document or the specification described therein, or any of its contents, either expressly or impliedly, or to any intellectual property owned or controlled by NVM Express, Inc., including, without limitation, any trademarks of NVM Express, Inc.

LEGAL DISCLAIMER:

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

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

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

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

NVM Express Technical Proposal for New Feature

Technical Proposal ID	4040a
Change Date	2020-08-25
Builds on Specification	NVM Express 1.4 TP 4056 Namespace Types

Technical Proposal Author(s)

Name	Company
David Black, Dmitry Tylik	Dell EMC

This technical proposal defines mechanisms to limit data size of commands that aren't subject to MDTs (Maximum Data Transfer Size) in Identify Controller – Dataset Management, Verify, Write Uncorrectable, and Write Zeroes.

Revision History

Revision Date	Change Description
2019-04-29	Initial version
2019-05-01	Revised text, added Data Set Management limits
2019-05-15	Minor edits, added notes about ONCS bit set vs. clear design issue.
2019-05-16	Revised based on comments from Mike Allison – important change is that size (number of logical blocks) limit on Dataset Management command field needs to be larger than one byte.
2019-05-17	Remove Editorial comments, move Write Uncorrectable Maximum Data Size (WUMDS) field to come before Dataset Management limit fields in order to make the four single-byte fields be adjacent. Reset change tracking
2019-05-24	Moved status code information from Identify Controller Data Structure to Dataset Management, Verify, Write Uncorrectable and Write Zeroes command specific sections.
2019-07-18	Minor edits, rebase to NVMe 1.4 from NVMe 1.3.
2019-07-30	Re-formatted changes to command descriptions based on comments from Mike Allison and Fred Knight. Consistently use “data size” instead of “amount of data”
2019-08-01	Prepare for phase 2 exit. One page is an acceptable data size limit for Write Uncorrectable and an acceptable range count limit for Dataset Management. Additional minor edits.
2019-09-26	Updated based on results of 9/19 straw poll. Controllers have a way to indicate partial support for Dataset Management, Verify, Write Uncorrectable and Write Zeroes commands. Optional NVM Command Support – Restricted (ONCSR) field was removed.
2019-10-01	Add additional text to Verify to complete ONCSR removal, plus notes to enable that design decision to be checked.

2019-10-05	Complete ONCSR removal for all commands and add text to use limit fields for recommended max sizes when commands are supported (i.e., when relevant ONCS bit is set to '1'). Remove new errors for Dataset Management command (it's still advisory, period) and add recommendations (should) that controller process portion of Dataset Management command up to any indicated limits. Serious editorial pass.
2019-10-08	Revise and clarify Dataset Management processing effects. Other minor edits.
2019-10-22	Minor edits
2019-10-24	Edits from Technical WG discussion.
2019-10-27	Edits including updates from on comments from John Meneghini and Christoph Hellwig. Use a new section # for the Dataset Management Limits, other minor edits. Retain field names that include the word "Limit" as hosts ought to treat these fields as limits in all cases, even though going beyond them does not always result in an error.
2019-11-12	Edits based on comments from Mike Allison.
2019-12-09	Add ONCS bit cleared to '0' case to Dataset Management command in order to match other commands. Add third Dataset Management limit (number of logical blocks in single range). Rearrange text so that the ONCS bit set to '1' case always comes first.
2019-12-10	Simplify Dataset Management limit functionality by requiring that either all three limit fields be non-zero or that all three limit fields be zero, as suggested by John Meneghini. Additional minor edits.
2019-12-12	Extensive edits to Dataset Management limits from meeting discussion: <ul style="list-style-type: none"> • Precisely specify effects of the three Dataset Management limits, including "shall not process" requirements when any limit is exceeded. • Add new Command Too Large status for limit exceeded errors (only occur when ONCS bit for Dataset Management is cleared to '0'). • Always use "all three" when referring to the three limits jointly. • Additional minor edits.
2019-12-13	In new Dataset Management processing limits text, add references to Figure 366 for range number and starting LBA of a range.
2019-12-17	Make it clear that Dataset Management limits are reported by the controller. Additional minor edits. Use "does not" instead of "shall not" for definitions of limit fields.
2019-12-19	For ONCS bit is set to '1' case, use "should not process" instead of "shall not process" for logical blocks in excess of limit. Rename Command Too Large status to Command Size Limit Exceeded, set off example in separate subsection. Add note warning that TP 4056 interaction may result in new limit fields moving into a new Identify data structure. Additional minor edits.
2019-12-30	Final edits for member review, based partly on comments from Mike Allison.
2020-01-17	Editorial changes based on member review comments from HPE. Update dates from 2019 to 2020.
2020-02-11	Post member-review: Remove comment on field relocation.
2020-02-13	Rewrite explanation of 255 max DMRL value (in new section 6.7.1) as a separate sentence. Remaining open issue is whether the new fields proposed for Identify Controller get moved to another data structure.
2020-03-04	Resolve open issue by moving new fields to new data structure defined in TP 4056 (Namespace Types). Refer to 2019-12-19 change description which indicates that the possibility of this move was indicated in the version of this TP that was sent to Member Review.
2020-06-07	Integrated into the NVM Express Base Specification
2020-06-11	Removing change bars.
2020-06-24	Updated definition of WUSL to reference Write Uncorrectable instead of Write Zeroes.

2020-06-26	Additional editorial cleanup
2020-08-18	As reported by Samsung during the 30 day review the definition of MDTS has an editorial update.
2020-08-20	Edit on new text to MDTS per Technical WG review. Accepted changes and removed comments.
2020-08-25	Integrated into the NVM Express Base Specification.

Discussion

This technical proposal addresses this issue:

The existing NVMe 1.4 specification enables the controller to specify Maximum Data Transfer Size for commands that transfer data. Currently there is no way for the controller to specify maximum data size for commands that do not transfer data such as Write Uncorrectable and Write Zeroes.

Description for NVMe Changes Document

- a) Specifying Size Limits for commands that do not transfer data:
 - a) Identify I/O Command Set specific Controller Data Structure for the NVM command set (defined in TP 4056) extended to include
 - o Dataset Management Ranges Limit (DMRL)
 - o Dataset Management Range Size Limit (DMRSL)
 - o Dataset Management Size Limit (DMSL)
 - o Verify Size Limit (VSL)
 - o Write Zeroes Size Limit (WZSL)
 - o Write Uncorrectable Size Limit (WUSL)

Description of Specification Changes

Modify Figure 129 (Status Code – Command Specific Status Values, NVM Command Set) as shown below:

Figure 1: Status Code – Command Specific Status Values, NVM Command Set

Value	Description	Commands Affected
80h	Conflicting Attributes	Dataset Management, Read, Write
81h	Invalid Protection Information	Compare, Read, Verify, Write, Write Zeroes
82h	Attempted Write to Read Only Range	Dataset Management, Write, Write Uncorrectable, Write Zeroes
83h	Command Size Limit Exceeded	Dataset Management
84h to BFh	Reserved	

Modify Portions of Figure 247 (Identify – Identify Controller Data Structure) as shown below:

Bytes	O/M ¹	Description
...
77	M	<p>Maximum Data Transfer Size (MDTS): This field indicates the maximum data transfer size for a command that transfers data between memory accessible by the host (e.g., host memory, Controller Memory Buffer (refer to Section 4.7 and the controller. The host should not submit a command that exceeds this maximum data transfer size. If a command is submitted that exceeds this transfer size, then the command is aborted with a status of Invalid Field in Command. The value is in units of the minimum memory page size (CAP.MPSMIN) and is reported as a power of two (2^n). A value of 0h indicates that there is no maximum data transfer size. This field includes the length of metadata, if metadata is interleaved with the logical block data. This field does not apply to commands that do not transfer data between memory accessible by the host and the controller (e.g., the Verify command, the Write Uncorrectable command, and the Write Zeroes command); there is no maximum data transfer size for these commands refer to the ONCS field for restrictions on these commands and other commands that transfer data.</p> <p>If SGL Bit Bucket descriptors are supported, their lengths shall be included in determining if a command exceeds the Maximum Data Transfer Size for destination data buffers. Their length in a source data buffer is not included for a Maximum Data Transfer Size calculation.</p>
...

Bytes	O/M ¹	Description
521:520	M	<p>Optional NVM Command Support (ONCS): This field indicates the optional NVM commands and features supported by the controller. Refer to section 6.</p> <p>Bits 15:8 are reserved.</p> <p>Bit 7 if set to '1', then the controller supports the Verify command and the Verify Size Limit (VSL) field indicates the recommended maximum data size for Verify commands. If cleared to '0', then the controller does not support of the Verify command is indicated by a non-zero data size limit in the VSL field.</p> <p>Bit 6 if set to '1', then the controller supports the Timestamp feature. If cleared to '0', then the controller does not support the Timestamp feature. Refer to section 5.21.1.14.</p> <p>Bit 5 if set to '1', then the controller supports reservations. If cleared to '0', then the controller does not support reservations. If the controller supports reservations then the following commands associated with reservations shall be supported: Reservation Report, Reservation Register, Reservation Acquire, and Reservation Release. Refer to section 8.8 for additional requirements.</p> <p>Bit 4 if set to '1', then the controller supports the Save field set to a non-zero value in the Set Features command and the Select field set to a non-zero value in the Get Features command. If cleared to '0', then the controller does not support the Save field set to a non-zero value in the Set Features command and the Select field set to a non-zero value in the Get Features command.</p> <p>Bit 3 if set to '1', then the controller supports the Write Zeroes command and the Write Zeroes Size Limit (WZSL) field indicates the recommended maximum data size for Write Zeroes commands. If cleared to '0', then the controller does not support of the Write Zeroes command is indicated by a non-zero data size limit in the WZSL field.</p> <p>Bit 2 if set to '1', then the controller supports the Dataset Management command and limits, if any, on controller support of the Dataset Management command are indicated by non-zero values in the Dataset Management Ranges Limit (DMRL) field, the Dataset Management Size Limit (DMSL) field and the Dataset Management Range Size Limit (DMRSL) field. If cleared to '0', then the controller does not support of the Dataset Management command is indicated by non-zero values in the DMRL, DMSL and DMRSL fields.</p> <p>Bit 1 if set to '1', then the controller supports the Write Uncorrectable command and the Write Uncorrectable Size Limit (WUSL) field indicates the recommended maximum data size for Write Uncorrectable commands. If cleared to '0', then the controller does not support of the Write Uncorrectable command is indicated by a non-zero data size limit in the WUSL field.</p> <p>Bit 0 if set to '1', then the controller supports the Compare command. If cleared to '0', then the controller does not support the Compare command.</p>
...

Modify the Identify Controller Data Structure, NVM Command Set defined in Section 5.15.2.TBDb.1 of TP 4056 (Namespace Types) by adding the following fields:

Bytes	O/M ¹	Description
...
00	O	<p>Verify Size Limit (VSL): If bit 7 in the Optional NVM Command Support (ONCS) field is set to '1' then:</p> <ul style="list-style-type: none"> a) a non-zero value in this field indicates the recommended maximum data size for a Verify command (refer to section 6.14); and b) a value of 0h in this field indicates that no recommended maximum data size for a Verify command is reported. <p>If bit 7 in the ONCS field is cleared to '0' then:</p> <ul style="list-style-type: none"> a) a non-zero value in this field indicates that the controller supports the Verify command with the maximum data size limit indicated by this field (refer to section 6.14); and b) a value of 0h in this field indicates that the controller does not support the Verify command. <p>The non-zero value is in units of the minimum memory page size (CAP.MPSMIN) and is reported as a power of two (2^n). This field includes the length of metadata, if metadata is interleaved with the logical block data.</p>
01	O	<p>Write Zeroes Size Limit (WZSL): If bit 3 in the Optional NVM Command Support (ONCS) field is set to '1' then:</p> <ul style="list-style-type: none"> a) a non-zero value in this field indicates the recommended maximum data size for a Write Zeroes command (refer to section 6.17); and b) a value of 0h in this field indicates that no recommended maximum data size for a Write Zeroes command is reported. <p>If bit 3 in the ONCS field is cleared to '0' then:</p> <ul style="list-style-type: none"> a) a non-zero value in this field indicates that the controller supports the Write Zeroes command with the maximum data size limit indicated by this field (refer to section 6.17); and b) a value of 0h in this field indicates that the controller does not support the Write Zeroes command. <p>The non-zero value is in units of the minimum memory page size (CAP.MPSMIN) and is reported as a power of two (2^n). This field includes the length of metadata, if metadata is interleaved with the logical block data.</p>

Bytes	O/M ¹	Description
02	O	<p>Write Uncorrectable Size Limit (WUSL): If bit 1 in the Optional NVM Command Support (ONCS) field is set to '1' then:</p> <ul style="list-style-type: none"> a) a non-zero value in this field indicates the recommended maximum data size for a Write Uncorrectable command (refer to section 6.16); and b) a value of 0h in this field indicates that no recommended maximum data size for a Write Uncorrectable command is reported. <p>If bit 1 in the ONCS field is cleared to '0' then:</p> <ul style="list-style-type: none"> a) a non-zero value in this field indicates that the controller supports the Write Uncorrectable command with the maximum data size limit indicated by this field (refer to section 6.16); and b) a value of 0h in this field indicates that the controller does not support the Write Uncorrectable command. <p>The non-zero value is in units of the minimum memory page size (CAP.MPSMIN) and is reported as a power of two (2^n). This field includes the length of metadata, if metadata is interleaved with the logical block data.</p>
03	O	<p>Dataset Management Ranges Limit (DMRL): If bit 2 in the Optional NVM Command Support (ONCS) field is set to '1' then:</p> <ul style="list-style-type: none"> a) a non-zero value in this field indicates the recommended maximum number of logical block ranges for a Dataset Management command (refer to section 6.7.1); and b) a value of 0h in this field indicates that no recommended maximum number of logical block ranges for a Dataset Management command is reported. <p>If bit 2 in the ONCS field is cleared to '0', then:</p> <ul style="list-style-type: none"> a) a non-zero value in this field indicates that the controller supports the Dataset Management command with the maximum number of logical block ranges limit indicated by this field (refer to section 6.7.1); and b) a value of 0h in this field indicates that the controller does not support the Dataset Management command.
07:04	O	<p>Dataset Management Range Size Limit (DMRSL): If bit 2 in the Optional NVM Command Support (ONCS) field is set to '1' then:</p> <ul style="list-style-type: none"> a) a non-zero value in this field indicates the recommended maximum number of logical blocks in a single range for a Dataset Management command (refer to section 6.7.1); and b) a value of 0h in this field indicates that no recommended maximum number of logical blocks in a single range for a Dataset Management command is reported. <p>If bit 2 in the ONCS field is cleared to '0', then:</p> <ul style="list-style-type: none"> a) a non-zero value in this field indicates that the controller supports the Dataset Management command with the maximum number of logical blocks in a single range limit indicated by this field (refer to section 6.7.1); and b) a value of 0h in this field indicates that the controller does not support the Dataset Management command.

Bytes	O/M ¹	Description
15:08	O	<p>Dataset Management Size Limit (DMSL): If bit 2 in the Optional NVM Command Support (ONCS) field is set to '1' then:</p> <ul style="list-style-type: none"> a) a non-zero value in this field indicates the recommended maximum total number of logical blocks for a Dataset Management command (refer to section 6.7.1). b) a value of 0h in this field indicates that no recommended maximum total number of logical blocks for a Dataset Management command is reported. <p>If bit 2 in the ONCS field is cleared to '0', then:</p> <ul style="list-style-type: none"> a) a non-zero value in this field indicates that the controller supports the Dataset Management command with the maximum total number of logical blocks limit indicated by this field (refer to section 6.7.1); and b) a value of 0h in this field indicates that the controller does not support the Dataset Management command.
...

Make changes to the following sections as shown below:

...

6.7 Dataset Management command

Insert the following text as a new Section 6.7.1 at the end of Section 6.7, renumbering the existing section 6.7.1 and subsequent sections as necessary:

6.7.1 Dataset Management Processing Limits

Processing limits for Dataset Management commands are indicated by non-zero values in three fields in the Identify Controller data structure (refer to Figure 247):

- a) A non-zero value in the Dataset Management Ranges Limit (DMRL) field indicates a processing limit on the number of ranges. If the controller reports a non-zero value in this field, then the controller does not process attributes and context attributes for any range whose range number (i.e., the number of the range as specified in the Range column of Figure 366) is greater than or equal to that non-zero value.
- b) A non-zero value in the Dataset Management Range Size Limit (DMRSL) field indicates a processing limit on the number of logical blocks in a single range. If the controller reports a non-zero value in this field, then the controller does not process attributes and context attributes for any logical block whose LBA offset from the starting LBA of the range (refer to Figure 366) that specifies the logical block is greater than or equal to that non-zero value.
- c) A non-zero value in the Dataset Management Size Limit (DMSL) field indicates a processing limit on the total number of logical blocks for the command. If the controller reports a non-zero value in this field, then the controller does not process attributes and context attributes for any logical block specified by any range for which the sum of:
 - a. the number of logical blocks, specified by lower numbered ranges, if any, otherwise zero; and
 - b. the LBA offset for that logical block from the starting LBA of that range,is greater than or equal to that non-zero value.

A logical block specified by a Dataset Management command satisfies all three of these processing limits if and only if each processing limit does not prevent controller processing of attributes and context attributes for that logical block. A logical block specified by a Dataset Management command does not satisfy a processing limit if that limit prevents controller processing of attributes and context attributes for that logical block.

The controller shall set all three processing limit fields (i.e., the DMRL, DMRSL and DMSL fields) to non-zero values or shall clear all three processing limit fields to 0h. A controller is able to impose a subset of the three processing limits by setting the field that reports each unused processing limit to the maximum possible value for that field (i.e., all bits set to '1'), with the exception that the resulting processing limit for the number of ranges is 255 of the 256 ranges supported by the Dataset Management command. Note that this exception is due to the DMRL field being 1-based in contrast to the 0's-based Number of Ranges (NR) field in the Dataset Management command.

If all three processing limit fields (i.e., the DMRL, DMRSL and DMSL fields) contain non-zero values, then the controller supports the Dataset Management command and:

- a) each processing limit field indicates a processing limit for controller processing of attributes and context attributes for logical blocks specified by the command;
- b) if bit 2 is set to '1' in the Optional NVM Command Support (ONCS) field in the Identify Controller data structure, then for the logical blocks specified by the command:

- a. the controller should process attributes and context attributes for all logical blocks that satisfy all three processing limits; and
 - b. the controller should not process attributes and context attributes for any logical blocks that do not satisfy one or more of the three processing limits;
- and
- c) if bit 2 is cleared to '0' in the ONCS field, then for the logical blocks specified by the command:
 - a. if all logical blocks specified by the command satisfy all three processing limits, then the controller shall process attributes and context attributes for those logical blocks; and
 - b. if the command specifies any logical block that does not satisfy one or more of the three processing limits, then the controller shall abort the command with Command Size Limit Exceeded status.

If all three processing limit fields (i.e., the DMRL, DMRSL and DMSL fields) are cleared to 0h then:

- a) if bit 2 is set to '1' in the ONCS field, then the controller supports the Dataset Management command and does not report any processing limits on the number of ranges, number of logical blocks in a single range or total number of logical blocks for the command; and
- b) if bit 2 is cleared to '0' in the ONCS field, then the controller does not support the Dataset Management command.

A controller may choose to take no action on any or all logical blocks for which attributes or context attributes are processed. If a Dataset Management command contains one or more ranges for which neither attributes nor context attributes are processed, then a controller may nonetheless check the fields that specify such ranges and abort the command if an error is detected (e.g., if the controller detects that such a range extends beyond the size of the namespace).

6.7.1.1 Dataset Management Processing Limits Example

For example, under the assumptions that the bit 2 in the ONCS field is set to '1' and the DMRSL field is set to its maximum value, consider a Dataset Management command that specifies two ranges, with range 0 containing 1024 logical blocks and range 1 containing 512 logical blocks:

- a) if the DMRL field is set to 1 and the DMSL field is set to 1,048, then the controller is expected to process attributes and context attributes for the logical blocks specified by range 0, and the controller does not process either attributes or context attributes for the logical blocks contained in range 1; and
- b) if the DMRL field is set to 2 and the DMSL field is set to 1,048, then the controller is expected to process attributes and context attributes for the logical blocks specified by range 0 and for the first 24 logical blocks of range 1, and the controller does not process either attributes or context attributes for the other logical blocks (i.e., 25 - 512) contained in range 1.

...

Modify Section 6.7.2 ([Dataset Management] Command Completion), which will be renumbered to Section 6.7.3, as follows

6.7.2 Command Completion

When the command is completed, the controller shall post a completion queue entry to the associated I/O Completion Queue indicating the status for the command.

Dataset Management command specific status values are defined in Figure 2.

Figure 2: Dataset Management – Command Specific Status Values

Value	Description
80h	Conflicting Attributes: The attributes specified in the command are conflicting.
82h	Attempted Write to Read Only Range: The controller may optionally report this status if a Deallocate is attempted for a read only range. The controller shall not return this status value if the read-only condition on the media is a result of a change in the write protection state of a namespace (refer to Section 368).
83h	Command Size Limit Exceeded: One or more of the Dataset Management processing limits (i.e., non-zero values of the DMRL, DMRSL and DMSL fields in the Identify Controller data structure) was exceeded (refer to section 6.7.1). The controller shall not return this status value if bit 2 is set to '1' in the Optional NVM Command Support field in the Identify Controller data structure.

6.14 Verify command

...

All data that is read or has its integrity checked by a Verify operation shall be included in the value of the Data Units Read field in the SMART/Health Information log page, refer to 5.14.1.2.

If the Verify Size Limit (VSL) field in the Identify Controller data structure is set to a non-zero value and:

- if bit 7 in the Optional NVM Command Support field in the Identify Controller data structure is set to '1' then the VSL field indicates the recommended maximum data size for the Verify command and any Verify command that specifies a logical block range whose data size exceeds that recommended maximum may encounter delays in processing; and
- if bit 7 in the ONCS field is cleared to '0' then the VSL field indicates the data size limit for the Verify command, and the controller shall abort any Verify command that specifies a logical block range whose data size exceeds that limit with a status of Invalid Field in Command.

...

6.16 Write Uncorrectable command

The Write Uncorrectable command is used to mark a range of logical blocks as invalid. When the specified logical block(s) are read after this operation, a failure is returned with Unrecovered Read Error status. To clear the invalid logical block status, a write operation is performed on those logical blocks.

If the Write Uncorrectable Size Limit (WUSL) field in the Identify Controller data structure is set to a non-zero value and:

- if bit 1 in the Optional NVM Command Support field in the Identify Controller data structure is set to '1' then the WUSL field indicates the recommended maximum data size for the Write Uncorrectable command and any Write Uncorrectable command that specifies a logical block range whose data size exceeds that recommended maximum may encounter delays in processing; and
- if bit 1 in the ONCS field is cleared to '0' then the WUSL field indicates the data size limit for the Write Uncorrectable command, and the controller shall abort any Write Uncorrectable command that specifies a logical block range whose data size exceeds that limit with a status of Invalid Field in Command.

...

6.17 Write Zeroes command

...

For each logical block in the range specified by a Write Zeroes command, if the namespace does not support that logical block clearing all bytes to 0h in the values read from that logical block and its metadata (excluding the protection information) read, then the controller shall not deallocate that logical block.

If the Write Zeroes Size Limit (WZSL) field in the Identify Controller data structure is set to a non-zero value, and:

- a) if bit 3 in the Optional NVMe Command Support field in the Identify Controller data structure is set to '1' then the WZSL field indicates the recommended maximum data size for the Write Zeroes command and any Write Zeroes command that specifies a logical block range whose data size exceeds that recommended maximum may encounter delays in processing; and
- b) if bit 3 in the ONCS field is cleared to '0' then the WZSL field indicates the data size limit for the Write Zeroes command, and the controller shall abort any Write Zeroes command that specifies a logical block range whose data size exceeds that limit with a status of Invalid Field in Command.

...