



LEGAL NOTICE:

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

This NVM Express Management Interface revision 1.1 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 Management Interface revision 1.1 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 - 2021 NVM Express, Inc. ALL RIGHTS RESERVED." When making any such citations or references to this document you are not permitted to revise, alter, modify, make any derivatives of, or otherwise amend the referenced portion of this document in any way without the prior express written permission of NVM Express, Inc. Nothing contained in this document shall be deemed as granting you any kind of license to implement or use this document or the specification described therein, or any of its contents, either expressly or impliedly, or to any intellectual property owned or controlled by NVM Express, Inc., including, without limitation, any trademarks of NVM Express, Inc.

LEGAL DISCLAIMER:

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

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

NVM Express Management Interface 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	6025a – Data Offset and Length Updates
Change Date	February 22, 2021
Builds on Specification	NVM Express Management Interface 1.1a

Technical Proposal Author(s)

Name	Company
Austin Bolen	Dell EMC

This TP deprecates the bits that indicate if Data Offset/Data Length are valid. The conditions under which these fields are valid are known a priori and already specified and so these bits are redundant. The bits will remain defined in order to interoperate with older Management Controllers but newer Management Endpoints will ignore them.

This TP also specifies when Data Offset/Data Length are within a valid range and specifies error status codes to report when they are not in range.

Revision History

Revision Date	Change Description
06/28/2020	Initial Draft
7/27/2020	Cleaned up language. Deleted old redundant text.
8/4/2020	Allow zero-length data transfers.
8/17/2020	Updates from Mike Allison and workgroup review on 8/10.
8/31/2020	Modified text to refer to implementations compliant to versions of this spec "later than 1.1" instead of "1.2 and later" to avoid referencing 1.2 which doesn't exist yet.
9/1/2020	Fixed a grammar issue.
10/12/2020	Editorial changes based on member review feedback.
10/14/2020	Integrated into the NVMe Management Interface Specification, Revision 1.1.
11/4/2020	Incorporated new phases for Invalid Parameter Error Responses.
11/23/2020	Moved DOFST/DLEN PEL changes from draft ECN003 content and merged them into this document.
11/24/2020	Fixed the Navigation page
12/7/2020	Added the Get Log Page command example. Removed the comments.
12/14/2020	Updated the filename to reflect in Phase 3.
12/17/2020	Accepted all changes and removed all comments for member review.
02/15/2021	Integrated into the NVMe Management Interface Specification, Revision 1.1.
02/22/2021	Accepted all changes, removed all comments, and changed all references/cross-references to text.

Description for NVMe Changes Document

Feature Enhancements:

- Data Offset and Data Length Updates (Mandatory)
 - Deprecate bits that indicate if Data Offset and Data Length fields are valid.
 - Specify additional cases where Data Offset/Data Length are within range or not.
 - Specify error status codes to report when Data Offset/Data Length are not in range.

Description of Specification Changes

Markup Conventions:

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

Modify Figure 112 (NVMe Admin Command Request Description) in section 6 as follows:

6 NVMe Express Admin Command Set

...

Figure 112: NVMe Admin Command Request Description

...	...								
05	<p>Command Flags (CFLGS): This field specifies flags for the command.</p> <p>Bits 7:2 are reserved.</p> <p>Bit 1, if set to '1', then the command contains an offset value in bytes 28-31. If cleared to '0', then the DOFST field shall be cleared to 0h.</p> <p>Bit 0, if set to '1', then the command contains a length value in bytes 32-35. If cleared to '0', then the DLEN field shall be cleared to 0h.</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>DOFST Valid (DOFSTV): This bit is not used and shall be ignored by the Management Endpoint for implementations compliant with versions of this specification later than 1.1.</td></tr> <tr> <td>0</td><td>DLEN Valid (DLENV): This bit is not used and shall be ignored by the Management Endpoint for implementations compliant with versions of this specification later than 1.1.</td></tr> </tbody> </table>	Bits	Description	7:2	Reserved	1	DOFST Valid (DOFSTV): This bit is not used and shall be ignored by the Management Endpoint for implementations compliant with versions of this specification later than 1.1.	0	DLEN Valid (DLENV): This bit is not used and shall be ignored by the Management Endpoint for implementations compliant with versions of this specification later than 1.1.
Bits	Description								
7:2	Reserved								
1	DOFST Valid (DOFSTV): This bit is not used and shall be ignored by the Management Endpoint for implementations compliant with versions of this specification later than 1.1.								
0	DLEN Valid (DLENV): This bit is not used and shall be ignored by the Management Endpoint for implementations compliant with versions of this specification later than 1.1.								
...	...								
31:28	<p>Data Offset (DOFST): For commands that transmit data from the Management Controller to the Management Endpoint (i.e., the Data Transfer subfield for the opcode field of the NVMe Admin Command as defined by the NVM Express Base Specification is 01b Request Data field in the Request Message has non-zero length) or do not transmit data (i.e., the Data Transfer subfield in the opcode field of the NVMe Admin Command as defined by the NVM Express Base Specification is 00b), this field shouldshall be cleared to 0h. If this field is not 0h for commands that transmit data from the Management Controller to the Management Endpoint or that do not transfer data, then the Management Endpoint shall return respond with an Invalid Parameter Error Response with the PEL field indicating this field.</p> <p>For commands that transmit data from the Management Endpoint to the Management Controller (i.e., the Response Data field in the Response Message has non-zero length), this field specifies the starting offset, in bytes, of the portion of data contained in the NVMe Admin Command completion data that is returned starting at byte offset 0h of the Response Data field in the Response Message.</p> <p>This field should be less than the size of the NVMe Admin Command completion data. If this field is greater than or equal to the size of the NVMe Admin Command completion data, then the Management Endpoint shall respond with an Invalid Parameter Error Response with the PEL field indicating this field.</p> <p>Bits 0 and 1 of this field shouldshall be cleared to '0'. If bits 0 and 1 are not cleared to '0', then the Management Endpoint should respond with an Invalid Parameter Error Response with the PEL field indicating this field.</p>								

Figure 112: NVMe Admin Command Request Description

35:32	<p>Data Length (DLEN): For commands that do not transfer data (i.e., the Data Transfer subfield in the opcode field of the NVMe Admin Command as defined by the NVM Express Base Specification is 00b), this field should be cleared to 0h. If this field is not 0h for commands that do not transfer data, then the Management Endpoint shall return respond with an Invalid Parameter Error Response with the PEL field indicating this field.</p> <p>For commands that transmit data from the Management Controller to the Management Endpoint (i.e., the Data Transfer subfield in the opcode field of the NVMe Admin Command as defined by the NVM Express Base Specification is 01b Request Data field in the Request Message has non-zero length), this field specifies the length, in bytes, of the data contained in the Request Data field in the Request Message.</p> <p>For commands that transmit data from the Management Endpoint to the Management Controller (i.e., the Data Transfer subfield in the opcode field of the NVMe Admin Command as defined by the NVM Express Base Specification is 10b Response Data field in the Response Message has non-zero length), this field specifies the length, in bytes, of the portion of data contained in the NVMe Admin Command completion data that is returned in the Response Data field in the Response Message. The sum of DLEN plus DOFST should be less than or equal to the size of the NVMe Admin Command completion data. If the sum is greater, then the Management Endpoint shall respond with an Invalid Parameter Error Response with the PEL indicating this field.</p> <p>For commands that transmit data (i.e., the Data Transfer subfield in the opcode of the NVMe Admin Command as defined by the NVM Express Base Specification is 01b or 10b), the Management Controller should specify a non-zero length in this field. If this field is cleared to 0h for commands that transmit data, then the Management Endpoint should respond with an Invalid Parameter Error Response with the PEL field indicating this field.</p> <p>Bits 0 and 1 of this field should be cleared to '0'. If bits 0 and 1 are not cleared to '0', then the Management Endpoint shall respond with an Invalid Parameter Error Response with the PEL field indicating this field.</p> <p>This field should be less than or equal to 4,096. If this field is greater than 4,096, then the Management Endpoint shall respond with an Invalid Parameter Error Response with the PEL field indicating this field.</p>
...	...

Modify a portion of Section 6.1 as follows:

6.1 Request and Response Data

NVMe Admin Commands may contain data as part of the Command Message. This data is passed in the Request Data field instead of using PRP Lists or SGL segments. The PRP Entry 2 (PRP2) and Metadata Pointer (MPTR) fields within the NVMe Admin Commands are reserved.

~~If there is no data sent with the NVMe Admin Command (e.g., the Data Transfer subfield for the opcode is 00b), then the Data Offset and Data Length fields shall be cleared to 0h.~~

~~If there is data sent with the NVMe Admin Command (i.e., the Data Transfer subfield for the opcode is 01b), then the Data Offset field shall be 0h and the Data Length field shall be set to the length of the Request Data required by the command. If the Data Length field does not correspond to the required length, the Management Endpoint shall respond with an Invalid Parameter Error Response.~~

If there is Response Data expected in the Response Message in the completion of the NVMe Admin Command (i.e., the Data Transfer subfield in the corresponding NVMe Admin Command for the opcode is 10b), then the Data Offset and Data Length fields describe the portion of the NVMe Admin Command completion data that is transferred in the Response Message. Any remaining data not transferred in the Response Message is discarded by the Management Endpoint as shown in Figure 115. ~~If the Data Length plus Data Offset fields are greater than the size of the NVMe Admin Command completion data, the Management Endpoint should respond with an Invalid Parameter Error Response.~~

Some NVMe Admin Commands specify an offset and length which shall be applied first to create the NVMe Admin Command completion data. Then DOFST and DLEN shall be applied which may further reduce the amount of Response Data. This results in Response Data starting from an offset into the NVMe Admin Command completion data equal to the sum of the offset specified by the NVMe Admin Command plus DOFST. Figure TBD provides an example for the Get Log Page command.

Figure 115: NVMe Admin Command Response Data Example

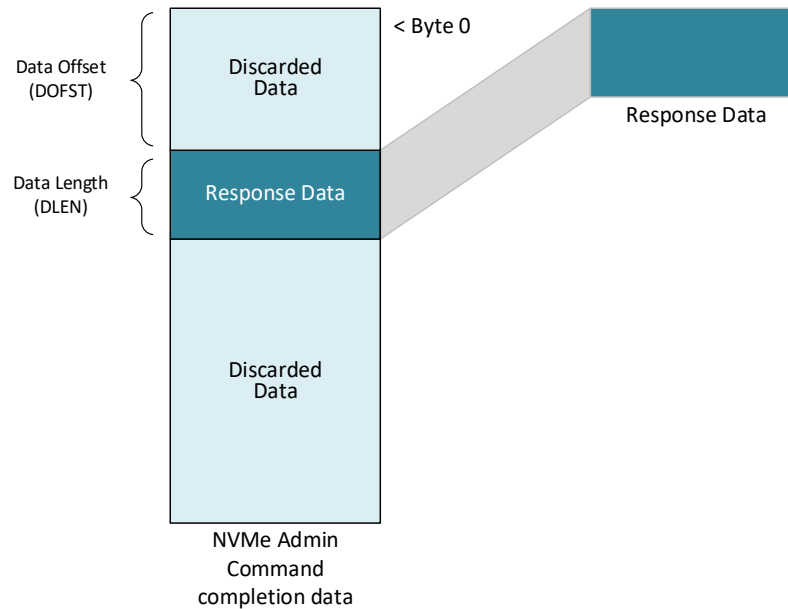


Figure TBD: NVMe Get Log Page Command Response Data Example

