



LEGAL NOTICE:

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

This NVM Express revision 2.0 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 2.0 technical proposal subject, however, to the Member's continued compliance with the Company's Intellectual Property Policy and Bylaws and the Member's Participation Agreement.

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

LEGAL DISCLAIMER:

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

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

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

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

NVM Express Technical Proposal for New Feature

Technical Proposal ID	4109 Allow host to specify Telemetry Host-Initiated data areas
Change Date	2021-07-16
Builds on Specification	NVM Express Base Specification 2.0
References Specification	

Technical Proposal Author(s)

Name	Company
Mike Allison, Andres Baez, Kapil Karkra	Intel
Lee Prewitt, Scott Lee	Microsoft

This proposal enhances the Get Log Page command to allow a host to specify the maximum data area to be created in the Telemetry Host-Initiated log page.

Revision History

Revision Date	Change Description
2021-02-15	Initial version
2021-02-17	Changes the MCDA field to be 3 bits and 1's based. The value of 000b specifies that the controller determines the maximum data area to be created in the log page. This makes the new field backwards compatible with existing host software that is not aware of the enhancement.
2021-03-24	Aligned to NVMe Base Specification 2.0 dated 3/23/2021. Using the LID Specific Field for specifying the support for the log page.
2021-04-13	Updated to the NVMe Base Specification 2.0 4/8/2021 version. Added the changes to allow each log page to define the use of the LID Specific Field.
2021-04-29	Accepted all changes and converted references/cross-references to text for member review.
2021-05-27	Added a table under MCDA that lists all possible values.
2021-06-03	Final wording on MCDA field. Approved during technical group meeting.
2021-06-14	Integrated into the NVMe Base Specification, revision 2.0.
2021-06-16	Addressed integration comments
2021-06-20	Integrated into the NVMe Base Specification, revision 2.0.
2021-07-16	Removed all comments, accepted all changes, and converted all references/cross-references to text.

Description for NVMe Base Specification 2.0 Changes Document

Feature Enhancement:

- **New requirement / incompatible change:**

Technical input submitted to the NVM Express™ Workgroup is subject to the terms of the NVM Express™ Participant's agreement. Copyright © 2014 to 2021 NVMe™ Corporation.

- A host is allowed to specify the maximum data area to be created in the Telemetry Host-Initiated log page in the Get Log Page commands.
- The LID Specific Field in the Supported Log Pages log page is now specified by each log page that utilizes the field.
- References:
 - NVM Express Base Specification 2.0 sections 15.16.1.8 and 5.16.1.14.

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 portions of NVM Express Base Specification 2.0 as shown below:

Modify a portion of 5.16.1.8 as shown below:

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 when 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 of Invalid Field in Command. This log page is global to the controller.

Figure 214: Command Dword 10 – Log Specific Field

Bits	Description
14:1209	Reserved
11:09	Maximum Created Data Area (MCDA): If the MCDAS bit is set to '1' in the LID Specific Field (refer to Figure FIG_ TBD) and the Create Telemetry Host-Initiated Data bit is set to '1', then this field specifies the data areas the host is requesting to be created in the log page. Data areas not requested shall not be created.
	If the MCDAS bit is cleared to '0' or the Create Telemetry Host-Initiated Data bit is cleared to '0', then this field is ignored.

Figure 214: Command Dword 10 – Log Specific Field

Bits	Description
08	Create Telemetry Host-Initiated Data: If set to '1', then the controller shall capture the Telemetry Host-Initiated Data representing the internal state of the controller at the time the associated Get Log Page command is processed. If cleared to '0', then the controller shall not update the Telemetry Host-Initiated Data. The Host-Initiated Data shall not change until the controller processes: <ul style="list-style-type: none">a) a subsequent Telemetry Host-Initiated Log with this bit set to '1';b) a Firmware Commit command; orc) a power on reset.

The Telemetry Host-Initiated Data consists of:

- a) Three areas, if bit 6 of the Log Page Attributes field is cleared to '0': Telemetry Host-Initiated Data Area 1, Telemetry Host-Initiated Data Area 2, and Telemetry Host-Initiated Data Area 3; or
- b) Four areas, if bit 6 of the Log Page Attributes field is set to '1': Telemetry Host-Initiated Data Area 1, Telemetry Host-Initiated Data Area 2, Telemetry Host-Initiated Data Area 3 and Telemetry Host-Initiated Data Area 4.

All areas start at Telemetry Host-Initiated Data Area Block 1. The last block of each area is indicated in Telemetry Host-Initiated Data Area y Last Block, respectively. The telemetry data captured and its size is implementation dependent.

The size of the log page is variable and:

- If bit 6 is cleared to '0' in the Log Page Attributes field, the size may be calculated using the Telemetry Host-Initiated Data Area 3 Last Block field.
- If bit 6 of the Log Page Attributes field is set to '1' and the Extended Telemetry Data Area 4 Supported (ETDAS) field is set to 1h in the Host Behavior Support feature (refer to section 5.21.1.22), then the size of the log page may be calculated using the Telemetry Host-Initiated Data Area 4 Last Block field.
- If bit 6 of the Log Page Attributes field is set to '1' and the Extended Telemetry Data Area 4 Supported (ETDAS) field is cleared to 0h in the Host Behavior Support feature (refer to section 5.21.1.22), then the size of the log page may be calculated using the Telemetry Host-Initiated Data Area 3 Last Block field.

The controller shall return data for all blocks requested:

- If bit 6 of the Log Page Attributes field is cleared to '0', then the data beyond the last block in Telemetry Host-Initiated Data Area 3 Last Block is undefined.
- If bit 6 of the Log Page Attributes field is set to '1' and the Extended Telemetry Data Area 4 Supported (ETDAS) field is set to 1h in the Host Behavior Support feature, then the data beyond the last block in Telemetry Host-Initiated Data Area 4 Last Block is undefined.
- If bit 6 of the Log Page Attributes field is set to '1' and the Extended Telemetry Data Area 4 Supported (ETDAS) field is cleared to 0h in the Host Behavior Support feature, then the data beyond the last block in Telemetry Host-Initiated Data Area 3 Last Block is undefined.

If the host requests a data transfer that is not a multiple of 512 bytes, then the controller shall return an error of Invalid Field in Command.

If the MCDAS bit (refer to [Figure FIG. TBD](#)) is set to '1' and a Get Log Page command with the Create Telemetry Host-Initiated Data bit is set to '1', then the maximum data area to be created in the Telemetry Host-Initiated log page shall be less than or equal to the MCDA field in the Log Specific Field field in Command Dword 10 (refer to [Figure 214](#)).

5.16.1.8.NEW Telemetry Host-Initiated LID Specific Field

[Figure FIG. TBD](#) specifies the format for the LID Specific Field field in the Supported Log Pages log page (refer to section [5.16.1.1](#)) for the Telemetry Host-Initiated log page.

Figure FIG_TBD: Telemetry Host-Initiated Log Page – LID Specific Field

Bits	Description
15:1	Reserved
0	Maximum Created Data Area Support (MCDAS): If set to '1', the controller supports the Maximum Created Data Area field in Log Specific Field field for a Telemetry Host-Initiated log page. If cleared to '0', then the controller does not support the Maximum Created Data Area field in Log Specific Field field for a Telemetry Host-Initiated log page.

Modify a portion of 5.16.1.1 as shown below:

5.16.1.1 Supported Log Pages (Log Identifier 00h)

An NVM subsystem may support several interfaces for submitting a Get Log Page command such as an Admin Submission Queue, PCIe VDM Management Endpoint, or SMBus/I2C Management Endpoint (refer to the NVM Express Management Interface Specification for details on Management Endpoints) and may have zero or more instances of each of those interfaces. The log pages supported on each instance of each interface may be different. This log page is used to describe the log pages that are supported on the interface to which the Get Log Page command was submitted and attributes specific to each log page. The log page is defined in Figure 203. The attributes of each log page are described in a LID Supported and Effects data structure defined in Figure 204.

The log pages that the controller supports are dependent on the I/O Command Set that is based on:

- the I/O Command Set selected in CC.CSS, if CC.CSS is not set to 110b; and
- the Command Set Identifier (CSI) field in CDW 14, if CC.CSS is set to 110b.

Figure 203: Supported Log Pages – Log

Bytes	Description
3:0	Log Page Identifier Supported 0: Contains the LID Supported and Effects data structure (refer to Figure 204.) for the LID 0h.
7:4	Log Page Identifier Supported 1: Contains the LID Supported and Effects data structure (refer to Figure 204.) for the LID 1h.
...	...
1019:1016	Log Page Identifier Supported 254: Contains the LID Supported and Effects data structure (refer to Figure 204.) for the LID FEh.
1023:1020	Log Page Identifier Supported 255: Contains the LID Supported and Effects data structure (refer to Figure 204.) for the LID FFh.

Figure 204: Get Log Page – LID Supported and Effects Data Structure

Bits	Description
31:16	LID Specific Field: This field is specific to the log page identifier and as defined in Figure 205. <i>if not defined for the log specified by the Log Page Identifier field, this field is reserved.</i>
31:2	Reserved
1	Index Offset Supported (IOS): If this bit is set to '1', then the controller supports an index offset for this LID in a Get Log Page command (i.e., the OT bit in the Get Log Page command is allowed to be set to '1'). If this bit is cleared to '0', then the controller does not support an index offset for this LID in a Get Log Page command (i.e., the OT bit in the Get Log Page command is only allowed to be cleared to '0').
0	LID Supported (LSUPP): If this bit is set to '1', then the controller supports this LID for a Get Log Page command. If this bit is cleared to '0', then the controller does not support this LID for a Get Log Page command. Refer to section 3.1.2 for the LID support requirements for each controller type.

Figure 205: LID Supported and Effects Data Structure – LID Specific Field

Log Page Identifier	LID Specific Field						
0 to Ch	Reserved						
0Dh	<p>The LID Specific Field for log page identifier 0Dh (Persistent Event Log) is defined as follows:</p> <table> <tr> <th>Bits</th><th>Description</th></tr> <tr> <td>15:1</td><td>Reserved</td></tr> <tr> <td>0</td><td> <p>Establish Context and Read 512 Bytes of Header Supported: If this bit is cleared to '0', then the controller does not support the Establish Context and Read 512 Bytes of Header action (refer to Figure 223).</p> <p>If this bit is set to '1', then the controller supports the Establish Context and Read 512 Bytes of Header action. If this bit is set to '1', then the Generation Number field in the Persistent Event Log shall also be supported.</p> <p>Implementations compliant to later than NVM Express revision 1.4 shall set this bit to '1'.</p> </td></tr> </table>	Bits	Description	15:1	Reserved	0	<p>Establish Context and Read 512 Bytes of Header Supported: If this bit is cleared to '0', then the controller does not support the Establish Context and Read 512 Bytes of Header action (refer to Figure 223).</p> <p>If this bit is set to '1', then the controller supports the Establish Context and Read 512 Bytes of Header action. If this bit is set to '1', then the Generation Number field in the Persistent Event Log shall also be supported.</p> <p>Implementations compliant to later than NVM Express revision 1.4 shall set this bit to '1'.</p>
Bits	Description						
15:1	Reserved						
0	<p>Establish Context and Read 512 Bytes of Header Supported: If this bit is cleared to '0', then the controller does not support the Establish Context and Read 512 Bytes of Header action (refer to Figure 223).</p> <p>If this bit is set to '1', then the controller supports the Establish Context and Read 512 Bytes of Header action. If this bit is set to '1', then the Generation Number field in the Persistent Event Log shall also be supported.</p> <p>Implementations compliant to later than NVM Express revision 1.4 shall set this bit to '1'.</p>						
0Eh to BFh	Reserved						
C0h to FFh	Vendor specific						

Modify a portion of 5.16.1.14 as shown below:

5.16.1.14 Persistent Event Log (Log Identifier 0Dh)

...

5.16.1.14.NEW Persistent Event Log LID Specific Field

Figure FIG_TBD specifies the format for the LID Specific Field field in the Supported Log Pages log page (refer to section 5.16.1.1) for the Persistent Event log page.

Figure FIG_TBD: Persistent Event Log Page – LID Specific Field

Bits	Description
15:1	Reserved
0	<p>Establish Context and Read 512 Bytes of Header Supported: If this bit is cleared to '0', then the controller does not support the Establish Context and Read 512 Bytes of Header action (refer to Figure 223).</p> <p>If this bit is set to '1', then the controller supports the Establish Context and Read 512 Bytes of Header action. If this bit is set to '1', then the Generation Number field in the Persistent Event Log shall also be supported.</p> <p>Implementations compliant to later than NVM Express revision 1.4 shall set this bit to '1'.</p>

5.16.1.14.1 Persistent Event Log Events

...