



**LEGAL NOTICE:**

© **Copyright 2007 – 2020 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 - 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.

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

<b>Technical Proposal ID</b>	<b>6017 – VPD Data Structure Enhancement</b>
<b>Change Date</b>	<b>June 3, 2020</b>
<b>Builds on Specification</b>	<b>NVM Express Management Interface 1.1</b>

### Technical Proposal Author(s)

Name	Company
Myron Loewen	Intel Corporation
Yoni Shternhell	WDC

The purpose of this Technical Proposal is to eliminate unintended fixed sizes in the VPD for the NVMe MultiRecord Area and the NVMe PCIe Port MultiRecord Area. These records have a length field that was set to a fixed value which required a total of 30 bytes of padding. By changing the length value to implementation specific as prior changes for NVMe-MI 1.1 made to the Product Info Area we achieve consistency across all the MultiRecord length fields. This offers implementers the flexibility of eliminating the padding when VPD approaches 256 bytes. Thus extending the life of designs before requiring implementation of TP6013. This TP does not require the elimination of the padding for products that have sufficient space.

**Revision History**

Revision Date	Change Description
02/03/2020	First published draft
02/10/2020	Added feedback from workgroup to make padding optional instead of variable sized
03/23/2020.	Cleaned up text change to standard if-then syntax and put reserved back in
03/24/2020	Fixed date, updated values to use xxh for consistency. Accepted all changes.
6/2/2020	Integrated into the NVMe Management Interface Specification, Revision 1.1.
6/3/2020	Accepted all changes
6/3/2020	Added showing change from 00h to 0h in figure 149.

## Description of Specification Changes

**Editor's note: These changes are applied on top of NVMe-MI 1.1.**

**Modify Figure 148 NVMe MultiRecord Area as shown below:**

Byte Offset	Factory Default	Description						
00	0Bh	NVMe Record Type ID						
01	02h or 82h	<b>Record Format:</b>						
		<table><tr><th>Bit</th><th>Definition</th></tr><tr><td>7</td><td>Set to '1' if last record in list.</td></tr><tr><td>6:0</td><td>Record format version = 2.</td></tr></table>	Bit	Definition	7	Set to '1' if last record in list.	6:0	Record format version = 2.
		Bit	Definition					
7	Set to '1' if last record in list.							
6:0	Record format version = 2.							
02	20h or 3Bh	<b>Record Length (RLEN):</b> This field indicates the length of the MultiRecord Area in bytes without including the first 5 bytes that are common to all MultiRecords.						
03	Impl Spec	<b>Record Checksum:</b> This field is used to give the record data a zero checksum (i.e., the modulo 256 sum of the record data bytes from byte offset 05 to the end of this record plus this checksum byte equals zero).						
04	Impl Spec	<b>Header Checksum:</b> This field is used to give the record header a zero checksum (i.e., the modulo 256 sum of the first byte of the header through this checksum byte equals zero).						
05	0h	<b>NVMe MultiRecord Area Version Number:</b> This field indicates the version number of this NVMe MultiRecord. This field shall be cleared to 0h in this version of the specification.						
06	Impl Spec	<b>Form Factor (FF):</b> This field indicates the form factor of the Management Endpoint. Refer to the values in Figure 157.						
12:07	0h	Reserved						
13	Impl Spec 1	<b>Initial 1.8 V Power Supply Requirements:</b> This field specifies the initial 1.8 V power supply requirements in Watts prior to receiving a Set Slot Power message.						
14	Impl Spec 1	<b>Maximum 1.8 V Power Supply Requirements:</b> This field specifies the maximum 1.8 V power supply requirements in Watts.						
15	Impl Spec 1	<b>Initial 3.3 V Power Supply Requirements:</b> This field specifies the initial 3.3 V power supply requirements in Watts prior to receiving a Set Slot Power message.						
16	Impl Spec 1	<b>Maximum 3.3 V Power Supply Requirements:</b> This field specifies the maximum 3.3 V power supply requirements in Watts.						
17	0h	Reserved						
18	Impl Spec 1	<b>Maximum 3.3 V aux Power Supply Requirements:</b> This field specifies the maximum 3.3 V power supply requirements in 10 mW units.						
19	Impl Spec 1	<b>Initial 5 V Power Supply Requirements:</b> This field specifies the initial 5 V power supply requirements in Watts prior to receiving a Set Slot Power message.						
20	Impl Spec 1	<b>Maximum 5 V Power Supply Requirements:</b> This field specifies the maximum 5 V power supply requirements in Watts.						
21	Impl Spec 1	<b>Initial 12 V Power Supply Requirements:</b> This field specifies the initial 12 V power supply requirements in Watts prior to receiving a Set Slot Power message.						
22	Impl Spec 1	<b>Maximum 12 V Power Supply Requirements:</b> This field specifies the maximum 12 V power supply requirements in Watts.						
23	Impl Spec	<b>Maximum Thermal Load:</b> This field specifies the maximum thermal load from the NVM Subsystem in Watts.						
36:24	Impl Spec	<b>Total NVM Capacity:</b> This field indicates the total NVM capacity of the NVM Subsystem in bytes.						

Byte Offset	Factory Default	Description
		If the NVM Subsystem supports Namespace Management, then this field should correspond to the value reported in the TNVMCAP field in the NVMe Identify Controller Data Structure.  A value of 0h may be used to indicate this feature is not supported.
63:37	0h	<b>Reserved</b> If the RLEN field is set to 3Bh, then this field is reserved. If the RLEN field is set to 20h, then this field is not present.
NOTES: 1. Power supply requirements shall be set to the smallest integer value which fully supplies the necessary power to the NVMe Storage Device. A value of 0h indicates that the power supply voltage is not used.		

**Modify Figure 149 NVMe PCIe Port MultiRecord Area as shown below:**

Byte Offset	Factory Default	Description												
00	0Ch	NVMe PCIe Port Record Type ID												
01	02h or 82h	<div>Record Format:<table><tr><th>Bit</th><th>Definition</th></tr><tr><td>7</td><td>Set to '1' if last record in list.</td></tr><tr><td>6:0</td><td>Record format version = 2.</td></tr></table></div>	Bit	Definition	7	Set to '1' if last record in list.	6:0	Record format version = 2.						
Bit	Definition													
7	Set to '1' if last record in list.													
6:0	Record format version = 2.													
02	08h or 0Bh	<div>Record Length (RLEN): This field indicates the length of the MultiRecord Area in bytes without including the first 5 bytes that are common to all MultiRecords.</div>												
03	Impl Spec	<div>Record Checksum: This field is used to give the record data a zero checksum (i.e., the modulo 256 sum of the record data bytes from byte offset 05 to the end of this record plus this checksum byte equals zero).</div>												
04	Impl Spec	<div>Header Checksum: This field is used to give the record header a zero checksum (i.e., the modulo 256 sum of the first byte of the header through this checksum byte equals zero).</div>												
05	1h	<div>NVMe PCIe Port MultiRecord Area Version Number: This field indicates the version number of this NVMe PCIe Port MultiRecord. This field shall be set to 1h in this version of the specification.</div>												
06	Impl Spec	<div>PCIe Port Number: This field contains the PCIe port number. This is the same value as that reported in the Port Number field in the PCIe Link Capabilities Register.</div>												
07	Impl Spec	<div>Port Information: This field indicates information about the PCIe Ports in the device.  Bits 7:1 are reserved.  Bit 0, if set to '1' indicates that all PCIe ports within the device have the same capabilities (i.e., the capabilities listed in this structure are consistent across each PCIe port).</div>												
08	Impl Spec	<div>PCIe Link Speed: This field indicates a bit vector of link speeds supported by the PCIe port.<table><tr><th>Bit</th><th>Definition</th></tr><tr><td>7:4</td><td>Reserved</td></tr><tr><td>3</td><td>Set to '1' if the PCIe link supports 16.0 GT/s, otherwise cleared to '0'.</td></tr><tr><td>2</td><td>Set to '1' if the PCIe link supports 8.0 GT/s, otherwise cleared to '0'.</td></tr><tr><td>1</td><td>Set to '1' if the PCIe link supports 5.0 GT/s, otherwise cleared to '0'.</td></tr><tr><td>0</td><td>Set to '1' if the PCIe link supports 2.5 GT/s, otherwise cleared to '0'.</td></tr></table></div>	Bit	Definition	7:4	Reserved	3	Set to '1' if the PCIe link supports 16.0 GT/s, otherwise cleared to '0'.	2	Set to '1' if the PCIe link supports 8.0 GT/s, otherwise cleared to '0'.	1	Set to '1' if the PCIe link supports 5.0 GT/s, otherwise cleared to '0'.	0	Set to '1' if the PCIe link supports 2.5 GT/s, otherwise cleared to '0'.
Bit	Definition													
7:4	Reserved													
3	Set to '1' if the PCIe link supports 16.0 GT/s, otherwise cleared to '0'.													
2	Set to '1' if the PCIe link supports 8.0 GT/s, otherwise cleared to '0'.													
1	Set to '1' if the PCIe link supports 5.0 GT/s, otherwise cleared to '0'.													
0	Set to '1' if the PCIe link supports 2.5 GT/s, otherwise cleared to '0'.													

Byte Offset	Factory Default	Description																														
09	Impl Spec	<p><b>PCIe Maximum Link Width:</b> The maximum PCIe link width for this NVM Subsystem port. This is the expected negotiated link width that the port link trains to if the platform supports it. A Requester may compare this value with the PCIe Negotiated Link Width to determine if there has been a PCIe link training issue.</p> <table><tr><th>Value</th><th>Definition</th></tr><tr><td>0</td><td>Reserved</td></tr><tr><td>1</td><td>PCIe x1</td></tr><tr><td>2</td><td>PCIe x2</td></tr><tr><td>3</td><td>Reserved</td></tr><tr><td>4</td><td>PCIe x4</td></tr><tr><td>5 to 7</td><td>Reserved</td></tr><tr><td>8</td><td>PCIe x8</td></tr><tr><td>9 to 11</td><td>Reserved</td></tr><tr><td>12</td><td>PCIe x12</td></tr><tr><td>13 to 15</td><td>Reserved</td></tr><tr><td>16</td><td>PCIe x16</td></tr><tr><td>17 to 31</td><td>Reserved</td></tr><tr><td>32</td><td>PCIe x32</td></tr><tr><td>33 to 255</td><td>Reserved</td></tr></table>	Value	Definition	0	Reserved	1	PCIe x1	2	PCIe x2	3	Reserved	4	PCIe x4	5 to 7	Reserved	8	PCIe x8	9 to 11	Reserved	12	PCIe x12	13 to 15	Reserved	16	PCIe x16	17 to 31	Reserved	32	PCIe x32	33 to 255	Reserved
Value	Definition																															
0	Reserved																															
1	PCIe x1																															
2	PCIe x2																															
3	Reserved																															
4	PCIe x4																															
5 to 7	Reserved																															
8	PCIe x8																															
9 to 11	Reserved																															
12	PCIe x12																															
13 to 15	Reserved																															
16	PCIe x16																															
17 to 31	Reserved																															
32	PCIe x32																															
33 to 255	Reserved																															
10	Impl Spec	<p><b>MCTP Support:</b> This field contains a bit vector that specifies the level of support for the NVMe Management Interface.</p> <p>Bits 7:1 are reserved.</p> <p>Bit 0, if set to '1' indicates that MCTP-based management commands are supported on the PCIe port.</p>																														
11	Impl Spec	<p><b>Ref Clk Capability:</b> This field contains a bit vector that specifies the PCIe clocking modes supported by the port.</p> <table><tr><th>Bit</th><th>Definition</th></tr><tr><td>7:4</td><td>Reserved</td></tr><tr><td>3</td><td>Set to '1' if the device automatically uses RefClk if provided and otherwise uses SRIS, otherwise cleared to '0'.</td></tr><tr><td>2</td><td>Set to '1' if the PCIe link supports Separate RefClk with SSC (SRIS), otherwise cleared to '0'.</td></tr><tr><td>1</td><td>Set to '1' if the PCIe link supports Separate RefClk with no SSC (SRNS), otherwise cleared to '0'.</td></tr><tr><td>0</td><td>Set to '1' if the PCIe link supports common RefClk, otherwise cleared to '0'.</td></tr></table>	Bit	Definition	7:4	Reserved	3	Set to '1' if the device automatically uses RefClk if provided and otherwise uses SRIS, otherwise cleared to '0'.	2	Set to '1' if the PCIe link supports Separate RefClk with SSC (SRIS), otherwise cleared to '0'.	1	Set to '1' if the PCIe link supports Separate RefClk with no SSC (SRNS), otherwise cleared to '0'.	0	Set to '1' if the PCIe link supports common RefClk, otherwise cleared to '0'.																		
Bit	Definition																															
7:4	Reserved																															
3	Set to '1' if the device automatically uses RefClk if provided and otherwise uses SRIS, otherwise cleared to '0'.																															
2	Set to '1' if the PCIe link supports Separate RefClk with SSC (SRIS), otherwise cleared to '0'.																															
1	Set to '1' if the PCIe link supports Separate RefClk with no SSC (SRNS), otherwise cleared to '0'.																															
0	Set to '1' if the PCIe link supports common RefClk, otherwise cleared to '0'.																															
12	Impl Spec	<p><b>Port Identifier:</b> This field contains the NVMe-MI Port Identifier.</p>																														
15:13	00h	<p><b>Reserved</b></p> <p>If the RLEN field is set to 0Bh, then this field is reserved. If the RLEN field is set to 08h, then this field is not present.</p>																														