



LEGAL NOTICE:

© **Copyright 2007 - 2018 NVM Express, Inc. ALL RIGHTS RESERVED.**

This NVM Express Management Interface revision 1.0a 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 1.0a 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 - 2018 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 NVM Express Administration
3855 SW 153rd Drive
Beaverton, OR 97003
admin@nvmexpress.org

NVM Express Technical Proposal for New Feature

Technical Proposal ID	6006 – PCIe Port Numbering
Change Date	1/28/2018
Builds on Specification	NVM Express Management Interface 1.0a

Technical Proposal Author(s)

Name	Company
Peter Onufryk	Microsemi

The NVMe-MI 1.0a specification requires that the Port Identifier for a PCIe port is the same as the Port Number field in the PCIe Link Capability Register. The effect of this requirement is that PCIe ports associated with NVM Subsystem need to have a unique PCIe port number. There is no requirement in the PCI Express Base Specification that ports associated with an Endpoint have unique port numbers. The purpose of this technical proposal is to remove the requirement that NVMe-MI imposes on PCIe.

Revision History

Revision Date	Change Description
10/9/2017	<ul style="list-style-type: none">Initial draft.
11/27/2017	<ul style="list-style-type: none">Updated text associated with Port Identifier field in the NVMe PCIe Port MultiRecord Area.
12/4/2017	<ul style="list-style-type: none">Fixed NVMe PCIe Port MultiRecord byte numbering.
1/28/2018	<ul style="list-style-type: none">Editorial changes to align with NVMe TP numbering, updated NVMe administration, and updated year.

Description of Specification Changes

Modify Section 1.4 as shown below:

1.4 Architectural Model

An NVMe storage device, such as a PCIe SSD, that implements this specification, consists of an NVM Subsystem with one or more Management Endpoints. There may be up to one Management Endpoint per PCIe port and SMBus/I2C port. Each Management Endpoint has a Port Identifier that is less than or equal to the Number of Ports (NUMP) field value in the NVM Subsystem Information Data Structure. ~~The Port Identifier for a PCIe port is the same as the Port Number field in the PCIe Link Capabilities Register.~~

NVMe-MI supports Vital Product Data (VPD) that utilizes the format defined in the IPMI Platform Management FRU Information Storage Definition and is stored in a FRU Information Device. The FRU Information Device may be implemented in the NVM Subsystem, in an external device (e.g., serial EEPROM), or a combination of the two. The VPD is accessible over any port that supports NVMe-MI using MCTP commands. If the NVMe storage device has an SMBus/I2C interface, then the VPD is accessible using the access mechanism over I2C as defined in the IPMI Platform Management FRU Information Storage Definition.

Modify Figure 64 as shown below:

Figure 64: PCIe Port Specific Data

Byte	Description																
08	<p>PCIe Maximum Payload Size: This field indicates the Max Payload Size for the specified PCIe port. If the link is not active, this field should be cleared to 0h.</p> <table><tr><th>Value</th><th>Definition</th></tr><tr><td>0h</td><td>128 bytes</td></tr><tr><td>1h</td><td>256 bytes</td></tr><tr><td>2h</td><td>512 bytes</td></tr><tr><td>3h</td><td>1024 bytes</td></tr><tr><td>4h</td><td>2048 bytes</td></tr><tr><td>5h</td><td>4096 bytes</td></tr><tr><td>6h-FFh</td><td>Reserved</td></tr></table>	Value	Definition	0h	128 bytes	1h	256 bytes	2h	512 bytes	3h	1024 bytes	4h	2048 bytes	5h	4096 bytes	6h-FFh	Reserved
Value	Definition																
0h	128 bytes																
1h	256 bytes																
2h	512 bytes																
3h	1024 bytes																
4h	2048 bytes																
5h	4096 bytes																
6h-FFh	Reserved																

09	<p>PCIe Supported Link Speeds Vector: This field indicates the Supported Link Speeds for the specified PCIe port.</p> <table> <tr> <th>Bit</th><th>Description</th></tr> <tr> <td>7:3</td><td>Reserved</td></tr> <tr> <td>2</td><td>This bit shall be set to '1' if the link supports 8.0 GT/s</td></tr> <tr> <td>1</td><td>This bit shall be set to '1' if the link supports 5.0 GT/s</td></tr> <tr> <td>0</td><td>This bit shall be set to '1' if the link supports 2.5 GT/s.</td></tr> </table>	Bit	Description	7:3	Reserved	2	This bit shall be set to '1' if the link supports 8.0 GT/s	1	This bit shall be set to '1' if the link supports 5.0 GT/s	0	This bit shall be set to '1' if the link supports 2.5 GT/s.																				
Bit	Description																														
7:3	Reserved																														
2	This bit shall be set to '1' if the link supports 8.0 GT/s																														
1	This bit shall be set to '1' if the link supports 5.0 GT/s																														
0	This bit shall be set to '1' if the link supports 2.5 GT/s.																														
10	<p>PCIe Current Link Speed: The port's PCIe negotiated link speed using the same encoding as the PCIe Supported Link Speed Vector field. A value of 0h in this field indicates the PCIe Link is not available.</p> <table> <tr> <th>Value</th><th>Definition</th></tr> <tr> <td>0h</td><td>Link not active</td></tr> <tr> <td>1h</td><td>The current link speed is the speed indicated in the supported link speed bit 0.</td></tr> <tr> <td>2h</td><td>The current link speed is the speed indicated in the supported link speed bit 1.</td></tr> <tr> <td>3h</td><td>The current link speed is the speed indicated in the supported link speed bit 2.</td></tr> <tr> <td>4h</td><td>The current link speed is the speed indicated in the supported link speed bit 3.</td></tr> <tr> <td>5h</td><td>The current link speed is the speed indicated in the supported link speed bit 4.</td></tr> <tr> <td>6h</td><td>The current link speed is the speed indicated in the supported link speed bit 5.</td></tr> <tr> <td>7h</td><td>The current link speed is the speed indicated in the supported link speed bit 6.</td></tr> <tr> <td>8h-FFh</td><td>Reserved</td></tr> </table>	Value	Definition	0h	Link not active	1h	The current link speed is the speed indicated in the supported link speed bit 0.	2h	The current link speed is the speed indicated in the supported link speed bit 1.	3h	The current link speed is the speed indicated in the supported link speed bit 2.	4h	The current link speed is the speed indicated in the supported link speed bit 3.	5h	The current link speed is the speed indicated in the supported link speed bit 4.	6h	The current link speed is the speed indicated in the supported link speed bit 5.	7h	The current link speed is the speed indicated in the supported link speed bit 6.	8h-FFh	Reserved										
Value	Definition																														
0h	Link not active																														
1h	The current link speed is the speed indicated in the supported link speed bit 0.																														
2h	The current link speed is the speed indicated in the supported link speed bit 1.																														
3h	The current link speed is the speed indicated in the supported link speed bit 2.																														
4h	The current link speed is the speed indicated in the supported link speed bit 3.																														
5h	The current link speed is the speed indicated in the supported link speed bit 4.																														
6h	The current link speed is the speed indicated in the supported link speed bit 5.																														
7h	The current link speed is the speed indicated in the supported link speed bit 6.																														
8h-FFh	Reserved																														
11	<p>PCIe Maximum Link Width: 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 Management Controller 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>5-7</td><td>Reserved</td></tr> <tr> <td>8</td><td>PCIe x8</td></tr> <tr> <td>9-11</td><td>Reserved</td></tr> <tr> <td>12</td><td>PCIe x12</td></tr> <tr> <td>13-15</td><td>Reserved</td></tr> <tr> <td>16</td><td>PCIe x16</td></tr> <tr> <td>17-31</td><td>Reserved</td></tr> <tr> <td>32</td><td>PCIe x32</td></tr> <tr> <td>33-255</td><td>Reserved</td></tr> </table>	Value	Definition	0	Reserved	1	PCIe x1	2	PCIe x2	3	Reserved	4	PCIe x4	5-7	Reserved	8	PCIe x8	9-11	Reserved	12	PCIe x12	13-15	Reserved	16	PCIe x16	17-31	Reserved	32	PCIe x32	33-255	Reserved
Value	Definition																														
0	Reserved																														
1	PCIe x1																														
2	PCIe x2																														
3	Reserved																														
4	PCIe x4																														
5-7	Reserved																														
8	PCIe x8																														
9-11	Reserved																														
12	PCIe x12																														
13-15	Reserved																														
16	PCIe x16																														
17-31	Reserved																														
32	PCIe x32																														
33-255	Reserved																														
12	<p>PCIe Negotiated Link Width: The negotiated PCIe link width for this port.</p> <table> <tr> <th>Value</th><th>Definition</th></tr> <tr> <td>0</td><td>Link not active</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>5-7</td><td>Reserved</td></tr> </table>	Value	Definition	0	Link not active	1	PCIe x1	2	PCIe x2	3	Reserved	4	PCIe x4	5-7	Reserved																
Value	Definition																														
0	Link not active																														
1	PCIe x1																														
2	PCIe x2																														
3	Reserved																														
4	PCIe x4																														
5-7	Reserved																														

		8	PCIe x8	
		9-11	Reserved	
		12	PCIe x12	
		13-15	Reserved	
		16	PCIe x16	
		17-31	Reserved	
		32	PCIe x32	
		33-255	Reserved	
13	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.			
31:14	Reserved			

Modify Section 9.2.14 as shown below:

9.2.14 NVMe PCIe Port MultiRecord Area

Byte Offset	Factory Default	Description						
00	0Ch	NVMe PCIe Port Record Type ID						
01	2h	Bit 7 – end of list; record format version = 2h						
02	28h	Record Length (RLEN): This field indicates the length of the MultiRecord Area in bytes.						
03	Impl Spec	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 through the end of this record plus this checksum byte equals zero)						
04	Impl Spec	Header Checksum: This field is used to give the record header a zero checksum (i.e., the modulo 256 sum of the preceding record bytes starting with the first byte of the header plus this checksum byte equals zero).						
05	01h	NVMe PCIe Port MultiRecord Area Version Number: This field indicates the version number of this multirecord. This field shall be set to zero one in this version of the specification.						
06	Impl Spec	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.						
07	Impl	Port Information: This field indicates information about the PCIe Ports in the device. Bits 7 to 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.						
08	Impl Spec	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>7:3</td><td>Reserved</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></table>	Bit	Definition	7:3	Reserved	2	Set to ‘1’ if the PCIe link supports 8.0 GT/s. Otherwise cleared to ‘0’.
Bit	Definition							
7:3	Reserved							
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'.																														
09	Impl Spec	<p>PCIe Maximum Link Width: 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 Management Controller 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>5-7</td><td>Reserved</td></tr><tr><td>8</td><td>PCIe x8</td></tr><tr><td>9-11</td><td>Reserved</td></tr><tr><td>12</td><td>PCIe x12</td></tr><tr><td>13-15</td><td>Reserved</td></tr><tr><td>16</td><td>PCIe x16</td></tr><tr><td>17-31</td><td>Reserved</td></tr><tr><td>32</td><td>PCIe x32</td></tr><tr><td>33-255</td><td>Reserved</td></tr></table>			Value	Definition	0	Reserved	1	PCIe x1	2	PCIe x2	3	Reserved	4	PCIe x4	5-7	Reserved	8	PCIe x8	9-11	Reserved	12	PCIe x12	13-15	Reserved	16	PCIe x16	17-31	Reserved	32	PCIe x32	33-255	Reserved
Value	Definition																																	
0	Reserved																																	
1	PCIe x1																																	
2	PCIe x2																																	
3	Reserved																																	
4	PCIe x4																																	
5-7	Reserved																																	
8	PCIe x8																																	
9-11	Reserved																																	
12	PCIe x12																																	
13-15	Reserved																																	
16	PCIe x16																																	
17-31	Reserved																																	
32	PCIe x32																																	
33-255	Reserved																																	
10	Impl Spec	<p>MCTP Support: This field contains a bit vector that specifies the level of support for the NVMe Management Interface.</p> <p>Bits 7 to 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>Ref Clk Capability: 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 ReClk with SSC (SRIS). Otherwise cleared to '0'.</td></tr><tr><td>1</td><td>Set to '1' if the PCIe link supports Separate ReClk with no SSC (SRNS). Otherwise cleared to '0'.</td></tr><tr><td>0</td><td>Set to '1' if the PCIe link supports common ReClk. 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 ReClk with SSC (SRIS). Otherwise cleared to '0'.	1	Set to '1' if the PCIe link supports Separate ReClk with no SSC (SRNS). Otherwise cleared to '0'.	0	Set to '1' if the PCIe link supports common ReClk. 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 ReClk with SSC (SRIS). Otherwise cleared to '0'.																																	
1	Set to '1' if the PCIe link supports Separate ReClk with no SSC (SRNS). Otherwise cleared to '0'.																																	
0	Set to '1' if the PCIe link supports common ReClk. Otherwise cleared to '0'.																																	
12	Impl Spec	Port Identifier: This field contains the NVMe-MI Port Identifier.																																
15:13	00h	Reserved																																

