



**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>6013 – FRU Information Device Extension</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
Austin Bolen	Dell EMC
Peter Onufryk	Intel Corporation

The purpose of this Technical Proposal is to specify the mechanisms and architecture necessary fit more content into the FRU Information Device. It is intended to be backwards compatible with legacy manageability systems but the extended content beyond 256B may not be accessible. It may be implemented in firmware to emulate a serial EEPROM but is not directly supported by typical off the shelf serial EEPROMs. In addition, some extra constraints are added with an implementation note for manageability systems to discover both 8- and 16-bit offset serial EEPROMs.

There are also changes to the number of write cycles remaining counter to make it more specific in an environment of devices with a variety of sizes beyond 256B. And a new Element for the Topology MultiRecord to indicate the max size of the VPD content stored in the FRU Information Device.

#### Revision History

Revision Date	Change Description
12/10/2019	First published draft based on prior power point discussions
2/21/2020	Updated to 2020 and ready for member review.
3/23/2020	Updated I2C Read figure and related text based on workgroup discussion
3/30/2020	More updates to I2C Read figure and test based on workgroup feedback
4/2/2020	Accepted all changes and removed comments for member review. Embedded the excel file as the diagram for future editing.
6/2/2020	Integrated into the NVMe Management Interface Specification, Revision 1.1.
6/3/2020	Updated to indicate figure 153 is modified from NVMe-MI 1.1 and not figure 154.

## Description of Specification Changes

**Editor's note: These changes are applied on top of NVMe-MI .next at the time of publishing. Note that yellow highlight is used to indicate all references that need hyperlinks. Figures and Tables with TBD values should be replaced with the correct sequential reference number in the final document.**

**Modify Figure 136 as shown below:**

Bytes	O/M <sup>1</sup>	Description								
252:240		Reserved								
253	M	<div>NVM Subsystem Report (NVMSR): This field reports information associated with the NVM Subsystem. At least one bit in this field shall be set to ‘1’.</div> <table><tr><th>Bits</th><th>Description</th></tr><tr><td>7:2</td><td>Reserved</td></tr><tr><td>1</td><td><b>NVMe Enclosure (NVMEE):</b> If set to ‘1’, then the NVM Subsystem is part of an NVMe Enclosure. If cleared to ‘0’, then the NVM Subsystem is not part of an NVMe Enclosure.</td></tr><tr><td>0</td><td><b>NVMe Storage Device (NVMESD):</b> If set to ‘1’, then the NVM Subsystem is part of an NVMe Storage Device. If cleared to ‘0’, then the NVM Subsystem is not part of an NVMe Storage Device.</td></tr></table>	Bits	Description	7:2	Reserved	1	<b>NVMe Enclosure (NVMEE):</b> If set to ‘1’, then the NVM Subsystem is part of an NVMe Enclosure. If cleared to ‘0’, then the NVM Subsystem is not part of an NVMe Enclosure.	0	<b>NVMe Storage Device (NVMESD):</b> If set to ‘1’, then the NVM Subsystem is part of an NVMe Storage Device. If cleared to ‘0’, then the NVM Subsystem is not part of an NVMe Storage Device.
Bits	Description									
7:2	Reserved									
1	<b>NVMe Enclosure (NVMEE):</b> If set to ‘1’, then the NVM Subsystem is part of an NVMe Enclosure. If cleared to ‘0’, then the NVM Subsystem is not part of an NVMe Enclosure.									
0	<b>NVMe Storage Device (NVMESD):</b> If set to ‘1’, then the NVM Subsystem is part of an NVMe Storage Device. If cleared to ‘0’, then the NVM Subsystem is not part of an NVMe Storage Device.									
254	M	<div><b>VPD Write Cycle Information (VWCI):</b> This field indicates information about remaining number of times that VPD contents are able to be updated using the VPD Write command.</div> <table><tr><th>Bits</th><th>Description</th></tr><tr><td>7</td><td><b>VPD Write Cycle Remaining Valid (VWCRV):</b> If this bit is set to ‘1’, then the VPD Write Cycle Remaining field is valid. If this bit is cleared to ‘0’, then the VPD Write Cycles Remaining field is invalid and cleared to ‘0’.</td></tr><tr><td>6:0</td><td><b>VPD Write Cycles Remaining (VWCR):</b> If the VPD Write Cycle Remaining Valid bit is set to ‘1’, then this field contains a value indicating the remaining number of times that VPD contents are able to be updated in units of 256 bytes using the VPD Write command. For example, a 1 KiB FRU Information Device that can be updated 8 times would indicate a value of 32 in this field. If this field is set to 7Fh, then the remaining number of times that VPD contents are able to be updated using the VPD Write command is greater than or equal to 7Fh. If the VPD Write Cycle Remaining Valid bit is cleared to ‘0’, then this field is not valid and shall be cleared to a value of 0h.</td></tr></table>	Bits	Description	7	<b>VPD Write Cycle Remaining Valid (VWCRV):</b> If this bit is set to ‘1’, then the VPD Write Cycle Remaining field is valid. If this bit is cleared to ‘0’, then the VPD Write Cycles Remaining field is invalid and cleared to ‘0’.	6:0	<b>VPD Write Cycles Remaining (VWCR):</b> If the VPD Write Cycle Remaining Valid bit is set to ‘1’, then this field contains a value indicating the remaining number of times that VPD contents are able to be updated in units of 256 bytes using the VPD Write command. For example, a 1 KiB FRU Information Device that can be updated 8 times would indicate a value of 32 in this field. If this field is set to 7Fh, then the remaining number of times that VPD contents are able to be updated using the VPD Write command is greater than or equal to 7Fh. If the VPD Write Cycle Remaining Valid bit is cleared to ‘0’, then this field is not valid and shall be cleared to a value of 0h.		
Bits	Description									
7	<b>VPD Write Cycle Remaining Valid (VWCRV):</b> If this bit is set to ‘1’, then the VPD Write Cycle Remaining field is valid. If this bit is cleared to ‘0’, then the VPD Write Cycles Remaining field is invalid and cleared to ‘0’.									
6:0	<b>VPD Write Cycles Remaining (VWCR):</b> If the VPD Write Cycle Remaining Valid bit is set to ‘1’, then this field contains a value indicating the remaining number of times that VPD contents are able to be updated in units of 256 bytes using the VPD Write command. For example, a 1 KiB FRU Information Device that can be updated 8 times would indicate a value of 32 in this field. If this field is set to 7Fh, then the remaining number of times that VPD contents are able to be updated using the VPD Write command is greater than or equal to 7Fh. If the VPD Write Cycle Remaining Valid bit is cleared to ‘0’, then this field is not valid and shall be cleared to a value of 0h.									
255	M	<div><b>Management Endpoint Capabilities (MEC):</b> This field indicates the capabilities of the Management Endpoint in the Controller.</div> <table><tr><th>Bits</th><th>Description</th></tr><tr><td>7:2</td><td>Reserved</td></tr><tr><td>1</td><td><b>PCIe Port Management Endpoint (PCIEME):</b> If set to ‘1’, then the NVM Subsystem contains a Management Endpoint on a PCIe port.</td></tr><tr><td>0</td><td><b>SMBus/I2C Port Management Endpoint (SMBUSME):</b> If set to ‘1’, then the NVM Subsystem contains a Management Endpoint on an SMBus/I2C port.</td></tr></table>	Bits	Description	7:2	Reserved	1	<b>PCIe Port Management Endpoint (PCIEME):</b> If set to ‘1’, then the NVM Subsystem contains a Management Endpoint on a PCIe port.	0	<b>SMBus/I2C Port Management Endpoint (SMBUSME):</b> If set to ‘1’, then the NVM Subsystem contains a Management Endpoint on an SMBus/I2C port.
Bits	Description									
7:2	Reserved									
1	<b>PCIe Port Management Endpoint (PCIEME):</b> If set to ‘1’, then the NVM Subsystem contains a Management Endpoint on a PCIe port.									
0	<b>SMBus/I2C Port Management Endpoint (SMBUSME):</b> If set to ‘1’, then the NVM Subsystem contains a Management Endpoint on an SMBus/I2C port.									

Bytes	O/M <sup>1</sup>	Description
NOTES:		
1. O/M definition: O = Optional, M = Mandatory.		


**Modify Section 9.2 as shown below:**

The Vital Product Data (VPD) is FRU information (refer to the IPMI Platform Management FRU Information Storage Specification) describing an NVMe Storage Device. Each NVMe Storage Device FRU shall have a FRU Information Device with a size of 256 to 4096 bytes which contains to hold the VPD as defined in the IPMI Platform Management FRU Information Storage Definition. The VPD for NVMe Storage Device FRUs shall contain the required elements defined in Figure 144. The VPD and FRU Information Device are optional for a) NVMe Storage Devices that are not FRUs (e.g., NVMe Storage Devices with a Form Factor type of Integrated per Figure 157), and b) NVMe Enclosures. The VPD contents for these optional use cases is outside the scope of this specification.

The VPD shall be accessible using the VPD Read command on all Management Endpoints on the NVMe Storage Device FRU. The entire contents of the VPD may be updated using the VPD Write command.

If the NVM Subsystem has an SMBus/I2C interface, then the VPD shall be accessible at the SMBus/I2C address of the FRU Information Device using the I2C Reads access mechanism over I2C as defined in the IPMI Platform Management FRU Information Storage Definition. Updating the VPD by writing to the FRU Information Device directly on SMBus/I2C using I2C Writes shall not be supported if the VPD Write command is supported. Refer to the IPMI Platform Management FRU Information Storage Definition for more information about the FRU Information Device access mechanisms (I2C Reads/I2C Writes) over SMBus/I2C.

**Figure TBD: I2C Read from a FRU Information Device**



The number of valid bits in the Command Offset is dependent on the Maximum FRU Information Size. Command Offset bits that contain an "X" in Figure TBD are not valid and shall be ignored. This example shows a FRU Information Device with 10 valid Command Offset bits which corresponds to a Maximum FRU Information Device Size of 1 KiB.

Figure TBD shows an I2C Read where the A6h addresses and Command Offset are provided by the Management Controller followed by data being returned from the Management Endpoint. The Command Offset as shown in Figure TBD is stored internal to the NVMe Storage Device (i.e., the internal offset).

If an I2C Read is issued, then data is returned from the internal offset within the FRU Information Device and then the internal offset is incremented by 1h. If the Management Controller reads the last byte of the FRU Information Device (refer to Maximum FRU Information Size) via an I2C Read, then the internal offset shall be cleared to 0h (i.e., rolls over to 0h). If only one byte of the Command Offset is provided by the Management Controller, then the least significant byte of the internal offset shall be set to that value and the most significant byte of the internal offset shall be cleared to 0h.

The internal offset shall be cleared to 0h following a power cycle of the FRU Information Device. Implementations are allowed to maintain the current internal offset value or clear it to 0h following a reset of the FRU Information Device.

**Modify Section 9.2.3 as shown below:**

This MultiRecord is used to describe the form factor, power requirements, and capacity of NVMe Storage Devices with a single NVM Subsystem. ~~Implementations compliant to version 1.1 and later of this specification should implement Starting with version 1.1 of this specification, this MultiRecord has been superseded by~~ the Topology MultiRecord (refer to section 9.2.5). For backwards compatibility ~~with Management Controllers designed to versions of this specification prior to 1.1~~, the NVMe MultiRecord and the NVMe PCIe Port MultiRecord (refer to section 9.2.4) should both be included in the VPD in addition to the Topology MultiRecord unless the NVMe Storage Device FRU has Expansion Connectors, has more than one NVM Subsystem, or if including both this MultiRecord and the NVMe PCIe Port MultiRecord would extend the size of the VPD beyond ~~the 256 bytes limit~~. If ~~either both~~ the NVMe MultiRecord ~~or and~~ NVMe PCIe Port MultiRecord ~~are not included cannot fit within the 256-byte size limit of the VPD~~ then neither MultiRecord should be included.

**Modify Section 9.2.4 as shown below:**

This MultiRecord is used to describe the PCIe connectivity for NVMe Storage Devices with a single NVM Subsystem. ~~Implementations compliant to version 1.1 and later of this specification should implement Starting with version 1.1 of this specification, this MultiRecord has been superseded by~~ the Topology MultiRecord (refer to section 9.2.5). For backwards compatibility ~~with Management Controllers designed to versions of this specification prior to 1.1~~, the NVMe PCIe Port MultiRecord and the NVMe MultiRecord (refer to section 9.2.3) should both be included in the VPD in addition to the Topology MultiRecord unless the NVMe Storage Device FRU has Expansion Connectors, has more than one NVM Subsystem, or if including both this MultiRecord and the NVMe MultiRecord would extend the size of the VPD beyond ~~the 256 bytes limit~~. If ~~either both~~ the NVMe MultiRecord ~~or and~~ NVMe PCIe Port MultiRecord ~~are not included cannot fit within the 256-byte size limit of the VPD~~ then neither MultiRecord should be included.

**Modify Figure 153 as shown below:**

Value	Name	Reference Section
0	Reserved	-
1	Extended Element Descriptor	9.2.5.1
2	Upstream Connector Element Descriptor	9.2.5.2
3	Expansion Connector Element Descriptor	9.2.5.3
4	Label Element Descriptor	9.2.5.4
5	SMBus/I2C Mux Element Descriptor	9.2.5.5
6	PCIe Switch Element Descriptor	9.2.5.6
7	NVM Subsystem Element Descriptor	9.2.5.7
<b>8</b>	<b>FRU Information Device Element Descriptor</b>	<b>9.2.5.TBD</b>
89 to 239	Reserved	-
240 to 255	Vendor specific	<b>9.2.5.8TBD</b>

**Insert new section 9.2.5.TBD before the prior Vendor Specific section to 9.2.5.8**

**9.2.5.TBD2 FRU Information Device Element Descriptor**

The FRU Information Device Element Descriptor is shown in Figure **TBD3** and is used to describe a FRU Information Device contained in the NVMe Storage Device.

Figure TBD3: FRU Information Device Element Descriptor

Byte Offset	Factory Default	Description																				
00	08h	<b>Type:</b> This field indicates the type of the Element Descriptor. The FRU Information Device Element Descriptor Type is 8.																				
01	00h	<b>Revision:</b> This field indicates the revision of the Element Descriptor. The FRU Information Device Element Descriptor Revision is 0h for this specification.																				
02	06h	<b>Length:</b> This field indicates the length of the FRU Information Device Element Descriptor in bytes.																				
03	A6h/A7h or 0h for NVM Storage Devices  A4h/A5h or 0h for Carriers	<b>SMBus/I2C Address Info:</b> If the NVMe Storage Device contains an SMBus/I2C port, then this field indicates the default SMBus/I2C addressing per the table below; else, this field shall be cleared to 0h. <table><tr><th>Bit</th><th>Description</th></tr><tr><td>7:1</td><td><b>SMBus/I2C Address:</b> This field contains the 7-bit SMBus/I2C address. Refer to Figure 15 for requirements.</td></tr><tr><td>0</td><td><b>ARP Capable:</b> If this bit is set to '1', then SMBus ARP is supported. If this bit is cleared to '0', then SMBus ARP is not supported. Refer to the SMBus Specification for additional details.</td></tr></table>	Bit	Description	7:1	<b>SMBus/I2C Address:</b> This field contains the 7-bit SMBus/I2C address. Refer to Figure 15 for requirements.	0	<b>ARP Capable:</b> If this bit is set to '1', then SMBus ARP is supported. If this bit is cleared to '0', then SMBus ARP is not supported. Refer to the SMBus Specification for additional details.														
Bit	Description																					
7:1	<b>SMBus/I2C Address:</b> This field contains the 7-bit SMBus/I2C address. Refer to Figure 15 for requirements.																					
0	<b>ARP Capable:</b> If this bit is set to '1', then SMBus ARP is supported. If this bit is cleared to '0', then SMBus ARP is not supported. Refer to the SMBus Specification for additional details.																					
04	Impl Spec	<b>SMBus/I2C Capabilities:</b> If the NVMe Storage Device contains an SMBus/I2C port, then this field indicates the SMBus/I2C capabilities per the table below; else, this field shall be cleared to 0h. <table><tr><th>Bit</th><th>Description</th></tr><tr><td>7</td><td><b>Reset:</b> If this bit is set to '1', then all of the SMBus/I2C reset mechanisms are supported as defined by the specification for the Form Factor in the Host Connector Element Descriptor.  If this bit is cleared to '0', then the FRU Information Device does not support all of the SMBus/I2C reset mechanisms defined by the specification for the Form Factor in the Host Connector Element Descriptor.</td></tr><tr><td>6</td><td><b>I2C Writes Allowed:</b> If this bit is set to '1', then the FRU Information Device is allowed to be written using an I2C Write operation.  If this bit is cleared to '0', then the FRU Information Device is not allowed to be written using an I2C Write operation.</td></tr><tr><td>5:2</td><td>Reserved</td></tr><tr><td>1:0</td><td><b>Maximum Speed:</b> This field is set to the highest supported SMBus/I2C clock speed supported by the FRU Information Device.<table><tr><th>Value</th><th>Description</th></tr><tr><td>0</td><td>100 kHz</td></tr><tr><td>1</td><td>400 kHz</td></tr><tr><td>2</td><td>1 MHz</td></tr><tr><td>3</td><td>Reserved</td></tr></table></td></tr></table>	Bit	Description	7	<b>Reset:</b> If this bit is set to '1', then all of the SMBus/I2C reset mechanisms are supported as defined by the specification for the Form Factor in the Host Connector Element Descriptor.  If this bit is cleared to '0', then the FRU Information Device does not support all of the SMBus/I2C reset mechanisms defined by the specification for the Form Factor in the Host Connector Element Descriptor.	6	<b>I2C Writes Allowed:</b> If this bit is set to '1', then the FRU Information Device is allowed to be written using an I2C Write operation.  If this bit is cleared to '0', then the FRU Information Device is not allowed to be written using an I2C Write operation.	5:2	Reserved	1:0	<b>Maximum Speed:</b> This field is set to the highest supported SMBus/I2C clock speed supported by the FRU Information Device. <table><tr><th>Value</th><th>Description</th></tr><tr><td>0</td><td>100 kHz</td></tr><tr><td>1</td><td>400 kHz</td></tr><tr><td>2</td><td>1 MHz</td></tr><tr><td>3</td><td>Reserved</td></tr></table>	Value	Description	0	100 kHz	1	400 kHz	2	1 MHz	3	Reserved
Bit	Description																					
7	<b>Reset:</b> If this bit is set to '1', then all of the SMBus/I2C reset mechanisms are supported as defined by the specification for the Form Factor in the Host Connector Element Descriptor.  If this bit is cleared to '0', then the FRU Information Device does not support all of the SMBus/I2C reset mechanisms defined by the specification for the Form Factor in the Host Connector Element Descriptor.																					
6	<b>I2C Writes Allowed:</b> If this bit is set to '1', then the FRU Information Device is allowed to be written using an I2C Write operation.  If this bit is cleared to '0', then the FRU Information Device is not allowed to be written using an I2C Write operation.																					
5:2	Reserved																					
1:0	<b>Maximum Speed:</b> This field is set to the highest supported SMBus/I2C clock speed supported by the FRU Information Device. <table><tr><th>Value</th><th>Description</th></tr><tr><td>0</td><td>100 kHz</td></tr><tr><td>1</td><td>400 kHz</td></tr><tr><td>2</td><td>1 MHz</td></tr><tr><td>3</td><td>Reserved</td></tr></table>	Value	Description	0	100 kHz	1	400 kHz	2	1 MHz	3	Reserved											
Value	Description																					
0	100 kHz																					
1	400 kHz																					
2	1 MHz																					
3	Reserved																					
05	8h to 0Ch inclusive	<b>Maximum FRU Information Device Size:</b> The maximum size of the FRU Information Device is 2 <sup>N</sup> bytes where N is the value in this field (e.g. a value of 8 in this field indicates a maximum FRU Information Device size of 2 <sup>8</sup> or 256 bytes).																				