



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	6022 – NVMe-MI 1.1 Maintenance
Change Date	2021.02.22
Builds on Specification	NVM Express Management Interface 1.1b NVM Express Base Specification 1.4b
References	TP 4071b CEL Enhancements TP 4052c Endurance Group Management TP 4056d Namespace Types TP 6026 Boot Partition Read Access via OOB TP 4046a Command Group Control Feature TP 6016a Port NVMe-MI section 8 to NVMe and extend Namespace Metadata

Technical Proposal Author(s)

Name	Company
Austin Bolen	Dell EMC
Myron Loewen	Intel
John Geldman	Kioxia

This proposal addresses the following:

- a) Updates for PCIe Gen 5 (link speeds and spec reference);
- b) Provide a mechanism to report which Feature IDs and Log Identifiers that are supported out-of-band; and
- c) Make the existing MsgTag example 6 in Appendix C normative that shows the MsgTag in a replayed response with RR0 > 0 is that of the replay request primitive (not that of the original message that is being replayed).

Revision History

Revision Date	Change Description
10/04/2020	<ul style="list-style-type: none">Initial Draft
10/26/2020	<ul style="list-style-type: none">Filled in the tables for which Log Identifiers and Feature Identifiers are mandatory, optional, or prohibited for SMBus/I2C and PCIe Management Endpoints on NVMe Storage Devices and NVMe Enclosures.
11/30/2020	<ul style="list-style-type: none">Added editor's note to change "http" to "https" in one of the spec URLs.Made I/O Command Set Profile optional for NVMe Storage Devices.Clarified that the supported FIDs/LIDs could vary between each instance of a given interface.
12/19/2020	<ul style="list-style-type: none">Noted that a few of the mandatory logs are optional for NVMe Enclosures compliant with versions 1.1 and earlier of this specification.Updated MCTP IDs and Codes (DSP0239) spec to 1.7.0 since that include PCIe 5.0 as a physical layer.Updated to note to editor to convert "http" to "https" for all URLs except for http://www.smbus.org as that is the only site in our references that does not support https.Updated text for changes document.Added more log pages from recent TPs and updated the list of reference documents to include those TPs.
02/20/2021	<ul style="list-style-type: none">Editorial updates based on member review comments.
02/22/2021	<ul style="list-style-type: none">Integrated into the NVMe Management Interface Specification, Revision 1.1 and the NVMe Base Specification.

Markup Conventions:

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

Description for NVMe-MI Changes Document

Feature Enhancements:

- PCIe Base Specification, revision 5.0 support
 - **Additional requirement / incompatible change** in sections **5.7**, **9.2.4**, **9.2.5.6**, and **9.5.2.7**:
 - Figures **92**, **149**, **166**, and **168** add a new enumeration for 32 GT/s link speed.
 - References:
 - NVMe Management Interface Specification, Revision 1.2 sections **5.7**, **9.2.4**, **9.2.5.6**, and **9.5.2.7**.
- Msg tag clarifications
 - **New requirement / incompatible change** in section **4.2.1.5**:
 - Added a requirement that the Msg tag in each packet of the replayed Response Message shall be set to the value of the Msg tag in the associated Replay Control Primitive.
 - References:
 - NVMe Management Interface Specification, Revision 1.2 section **4.2.1.5**.
- Log Identifiers supported
 - **New requirement / incompatible change** in section **6.TBD**:
 - Added Figures **TBD1** and **TBD2** to specify which Log Identifiers are mandatory, optional, or required for each type of Management Endpoint.
 - References:
 - NVMe Management Interface Specification, Revision 1.2 section **6.TBD**.
- Feature Identifiers supported
 - **New requirement / incompatible change** in section **6.4**:
 - Added Figures **TBD3** and **TBD4** to specify which Feature Identifiers are mandatory, optional, or required for each type of Management Endpoint.
 - References:
 - NVMe Management Interface Specification, Revision 1.2 section **6.4**.

Description for NVMe Base Spec Changes Document

Feature Enhancements:

<Note to editor: these changes modify TP 4071b so should be listed with those changes>

- Log Identifier and Feature Identifier support
 - **New requirement / incompatible change** in section 5.14.**TBD2** and 5.14.**TBD3**:
 - Added a requirement that the list of Log Identifiers returned in Supported Log Pages (Log Identifier 0h) is specific to the interface on which the Get Log Page command was submitted.
 - Added a requirement that the list of Feature Identifiers returned in Feature Identifiers Supported and Effects (Log Identifier 12h) is specific to the interface on which the Get Log Page command was submitted.
 - References:
 - NVMe Base Specification, Revision 2.0 section 5.14.**TBD2** and 5.14.**TBD3**.

Description of Specification Changes

Modify portions of the NVMe Management Interface Spec 1.1b as shown below:

Modify a portion of Section 1.11 (References) as follows:

<Note to editor: Add an “s” to the end of “http” in all URLs in section 1.11 except for <http://www.smbus.org>>

PCI Express Base Specification, revision [45.0](https://www.pcisig.com). Available from <https://www.pcisig.com>.

MCTP IDs and Codes (DSP0239), version 1.[57.0](https://dmtd.org). Available from <https://dmtd.org>.

Modify a portion of Section 4.2.1.5 (Replay) as follows:

The Replay Control Primitive is used to retransmit the Response Message for the last Command Message processed in a Command Slot. The replayed Response Message forms a new MCTP Response Message with Message Data starting from Response Replay Offset of the original Response Message and continuing to the end of the Response Message, including the original MIC. The first packet shall have SOM set to '1' and shall include the Message Header of the original Response Message even if the Response Replay Offset is not zero. [The Msg tag in each packet of the replayed Response Message shall be set to the value of the Msg tag in the associated Replay Control Primitive. Refer to the MCTP Base Specification for the definition of the Msg tag.](#)

Modify a portion of Figure 92 (PCIe Port Specific Data) as follows:

09	PCIe Supported Link Speeds Vector: This field indicates the Supported Link Speeds for the specified PCIe port.	
	Bit	Description
	7:54	Reserved
	4	Set to '1' if the PCIe link supports 32.0 GT/s, otherwise cleared to '0'.
	3	This bit shall be set to '1' if the link supports 16.0 GT/s. Set to '1' if the PCIe link supports 16.0 GT/s, otherwise cleared to '0'.
	2	This bit shall be set to '1' if the link supports 8.0 GT/s. Set to '1' if the PCIe link supports 8.0 GT/s, otherwise cleared to '0'.
	1	This bit shall be set to '1' if the link supports 5.0 GT/s. Set to '1' if the PCIe link supports 5.0 GT/s, otherwise cleared to '0'.
	0	This bit shall be set to '1' if the link supports 2.5 GT/s. Set to '1' if the PCIe link supports 2.5 GT/s, otherwise cleared to '0'.

Modify a portion of Figure 149 (NVMe PCIe Port MultiRecord Area) as follows:

08	Impl Spec	PCIe Link Speed: This field indicates a bit vector of link speeds supported by the PCIe port.	
		Bit	Definition
		7:54	Reserved
		4	Set to '1' if the PCIe link supports 32.0 GT/s, otherwise cleared to '0'.
		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'.

Modify a portion of Figure 166 (PCIe Switch Port Descriptor) as follows:

02	Impl Spec	PCIe Link Speed: This field indicates a bit vector of link speeds supported by the PCIe port.	
		Bit	Description
		7: 5 4	Reserved
		4	Set to '1' if the PCIe link supports 32.0 GT/s, otherwise cleared to '0'.
		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'.

Modify a portion of Figure 168 (NVM Subsystem Port Descriptor) as follows:

02	Impl Spec	PCIe Link Speed: This field indicates a bit vector of link speeds supported by the PCIe port.	
		Bit	Description
		7: 5 4	Reserved
		4	Set to '1' if the PCIe link supports 32.0 GT/s, otherwise cleared to '0'.
		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'.

Add the following to a new section after Section 6.2 (Status):

Section 6.TBD Get Log Page

Figure TBD1 and Figure TBD2 define the log pages that are mandatory, optional, and prohibited for SMBus/I2C and PCIe VDM Management Endpoints on NVMe Storage Devices and NVMe Enclosures.

<note to editor: In Figure TBD1, Media Unit Status and Supported Capacity Configuration List are defined in TP 4052c Endurance Group Management. Boot Partition is defined in TP 6026 Boot Partition Read Access via OOB. Command and Feature Lockdown is defined in TP 4046a Command Group Control Feature.>

Figure TBD1: Management Endpoint – Log Page Support

<u>Log Page Name</u>	<u>SMBus/I2C Log Page Support Requirements¹</u>		<u>PCIe VDM Log Page Support Requirements¹</u>	
	<u>NVMe Storage Device</u>	<u>NVMe Enclosure</u>	<u>NVMe Storage Device</u>	<u>NVMe Enclosure</u>
<u>Supported Log Pages</u>	<u>M²</u>	<u>M²</u>	<u>M²</u>	<u>M²</u>
<u>Error Information</u>	<u>M</u>	<u>M</u>	<u>M</u>	<u>M</u>
<u>SMART / Health Information (Controller scope)</u>	<u>M</u>	<u>O</u>	<u>M</u>	<u>O</u>
<u>SMART / Health Information (NVM subsystem scope)</u>	<u>O</u>	<u>O</u>	<u>O</u>	<u>O</u>
<u>Firmware Slot Information</u>	<u>M</u>	<u>O</u>	<u>M</u>	<u>O</u>
<u>Changed Namespace List</u>	<u>O</u>	<u>O</u>	<u>O</u>	<u>O</u>
<u>Commands Supported and Effects</u>	<u>O</u>	<u>O</u>	<u>O</u>	<u>O</u>
<u>Device Self-test</u>	<u>O</u>	<u>O</u>	<u>O</u>	<u>O</u>
<u>Telemetry Host-Initiated</u>	<u>O</u>	<u>O</u>	<u>O</u>	<u>O</u>
<u>Telemetry Controller-Initiated</u>	<u>O</u>	<u>O</u>	<u>O</u>	<u>O</u>
<u>Endurance Group Information</u>	<u>O</u>	<u>O</u>	<u>O</u>	<u>O</u>
<u>Predictable Latency Per NVM Set</u>	<u>O</u>	<u>O</u>	<u>O</u>	<u>O</u>
<u>Predictable Latency Event Aggregate</u>	<u>O</u>	<u>O</u>	<u>O</u>	<u>O</u>
<u>Asymmetric Namespace Access</u>	<u>O</u>	<u>O</u>	<u>O</u>	<u>O</u>
<u>Persistent Event</u>	<u>O</u>	<u>O</u>	<u>O</u>	<u>O</u>
<u>LBA Status Information</u>	<u>O</u>	<u>O</u>	<u>O</u>	<u>O</u>
<u>Endurance Group Event Aggregate</u>	<u>O</u>	<u>O</u>	<u>O</u>	<u>O</u>
<u>Media Unit Status</u>	<u>O</u>	<u>O</u>	<u>O</u>	<u>O</u>
<u>Supported Capacity Configuration List</u>	<u>O</u>	<u>O</u>	<u>O</u>	<u>O</u>
<u>Feature Identifiers Supported and Effects</u>	<u>M²</u>	<u>O</u>	<u>M²</u>	<u>O</u>
<u>NVMe-MI Commands Supported and Effects</u>	<u>O</u>	<u>O</u>	<u>O</u>	<u>O</u>
<u>Command and Feature Lockdown</u>	<u>O</u>	<u>O</u>	<u>O</u>	<u>O</u>
<u>Boot Partition</u>	<u>O</u>	<u>O</u>	<u>O</u>	<u>O</u>
<u>Notes:</u> 1. <u>O = Optional, M = Mandatory, P = Prohibited.</u> 2. <u>Optional for versions 1.1 and earlier of this specification.</u>				

Figure TBD2: Management Endpoint – NVM Command Set Specific Log Page Support

<u>Log Page Name</u>	<u>SMBus/I2C Log Page Support Requirements</u> ¹		<u>PCIe VDM Log Page Support Requirements</u> ¹	
	<u>NVMe Storage Device</u>	<u>NVMe Enclosure</u>	<u>NVMe Storage Device</u>	<u>NVMe Enclosure</u>
<u>Reservation Notification</u>	<u>O</u>	<u>O</u>	<u>O</u>	<u>O</u>
<u>Sanitize Status</u>	<u>O</u>	<u>O</u>	<u>O</u>	<u>O</u>
<u>Notes:</u> 1. <u>O = Optional, M = Mandatory, P = Prohibited.</u>				

Add the following to the end of Section 6.4 (Set Features and Get Features):

Figure [TBD3](#) and Figure [TBD4](#) define the features that are mandatory, optional, and prohibited for SMBus/I2C and PCIe VDM Management Endpoints on NVMe Storage Devices and NVMe Enclosures.

<note to editor: In Figure [TBD3](#), I/O Command Set Profile is defined in TP 4056d Namespace Types. Enhanced Controller Metadata, Controller Metadata, and Namespace Metadata are defined in TP 6016a.>

Figure [TBD3](#): Management Endpoint – Feature Support

<u>Feature Name</u>	<u>SMBus/I2C Log Page</u> <u>Support Requirements</u> ¹		<u>PCIe VDM Log Page</u> <u>Support Requirements</u> ¹	
	<u>NVMe</u> <u>Storage</u> <u>Device</u>	<u>NVMe</u> <u>Enclosure</u>	<u>NVMe</u> <u>Storage</u> <u>Device</u>	<u>NVMe</u> <u>Enclosure</u>
Arbitration	P	P	P	P
Power Management	O	O	O	O
LBA Range Type	P	P	P	P
Temperature Threshold	O	O	O	O
Error Recovery	P	P	P	P
Volatile Write Cache	P	P	P	P
Number of Queues	P	P	P	P
Interrupt Coalescing	P	P	P	P
Interrupt Vector Configuration	P	P	P	P
Write Atomicity Normal	P	P	P	P
Asynchronous Event Configuration	P	P	P	P
Autonomous Power State Transition	O	O	O	O
Host Memory Buffer	P	P	P	P
Timestamp	O	O	O	O
Keep Alive Timer	P	P	P	P
Host Controlled Thermal Management	O	O	O	O
Non-Operational Power State Config	O	O	O	O
Read Recovery Level Config	P	P	P	P
Predictable Latency Mode Config	P	P	P	P
Predictable Latency Mode Window	P	P	P	P
LBA Status Information Attributes	P	P	P	P
Host Behavior Support	P	P	P	P
Sanitize Config	O	O	O	O
Endurance Group Event Configuration	P	P	P	P
I/O Command Set Profile	O	P	O	P
Enhanced Controller Metadata	O	O	O	O
Controller Metadata	O	O	O	O
Namespace Metadata	O	O	O	O
<u>Notes:</u>				
1. O = Optional, M = Mandatory, P = Prohibited for Set Features/Optional for Get Features.				

Figure TBD4: Management Endpoint – NVM Command Set Specific Feature Support

<u>Feature Name</u>	<u>SMBus/I2C Log Page Support Requirements</u> ¹		<u>PCIe VDM Log Page Support Requirements</u> ¹	
	<u>NVMe Storage Device</u>	<u>NVMe Enclosure</u>	<u>NVMe Storage Device</u>	<u>NVMe Enclosure</u>
<u>Software Progress Marker</u>	<u>P</u>	<u>P</u>	<u>P</u>	<u>P</u>
<u>Host Identifier</u>	<u>P</u>	<u>P</u>	<u>P</u>	<u>P</u>
<u>Reservation Notification Mask</u>	<u>P</u>	<u>P</u>	<u>P</u>	<u>P</u>
<u>Reservation Persistence</u>	<u>P</u>	<u>P</u>	<u>P</u>	<u>P</u>
<u>Namespace Write Protection Config</u>	<u>P</u>	<u>P</u>	<u>P</u>	<u>P</u>
<u>Notes:</u> 1. <u>O = Optional, M = Mandatory, P = Prohibited for Set Features/Optional for Get Features.</u>				

Modify portions of TP 4071b as shown below:

Modify a portion of Section 5.14.TBD2 Supported Log Pages (Log Identifier 00h) as follows:

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 ~~the controller supports~~ and attributes specific to each log page. The log page is defined in Figure <TBD_FX1>. The attributes of each log page are described in a LID Supported and Effects data structure defined in Figure <TBD_FX2>.

Modify a portion of 5.14.TBD3 Feature Identifiers Supported and Effects (Log Identifier 12h) as follows:

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 feature identifiers (FIDs) supported on each instance of each interface may be different. This log page describes the ~~feature identifiers (FIDs)~~ that are supported on the interface to which the Get Log Page command was submitted ~~the controller supports~~ and the effects of those features on the state of the NVM subsystem. The log page is defined in Figure <TBD_FY1>. Each Feature Identifier's effects are described in a FID Supported and Effects data structure defined in Figure <TBD_FY2>.