



LEGAL NOTICE:

© Copyright 2007 to 2020 NVM Express™, Inc. ALL RIGHTS RESERVED.

This erratum to the NVM Express Management Interface revision 1.1 specification 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 erratum to the NVM Express Management Interface revision 1.1 specification 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 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.

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 Errata

Errata ID	001
Revision Date	2020-04-20
Affected Spec Ver.	NVM Express Management Interface 1.1
Corrected Spec Ver.	

Errata Author(s)

Name	Company
Austin Bolen	Dell EMC

Errata Overview

- Corrected the size of the Element Type field in Element Metadata Structure to be 5 bits instead of 6.
- Clarification to use proper field name for Data Length in VPD Write section.
- Grammar fix in Section 1.4 (NVM Subsystem Architectural Model).
- Require specifying Controller ID in Read NVMe-MI Data Structure reading a Data Structure Type (DTYPE) of 04h (Optional Commands Supported) since the optional commands can vary per controller.
- Corrected name of Data Structure Type 04h from “Optional Commands Supported” to “Controller Optional Commands Supported”.
- Added a missing bit (CSTS) to the list of error selection bits in two places in Report All.
- Removed redundant statement in 9.3.3 that Management Endpoints are reset by PCIe Link Down since that is already covered by PCI Express conventional reset.
- Removed duplicate sentence in Invalid Parameter status code definition.
- Added clarification to 5.1.3 that changing MTU with any outstanding Response Messages results in undefined behavior.
- Added clarification to 5.1.3 that changing MTU after sending a Command Message but before a Replay results in undefined behavior.
- Clarified that a Command Slot is paused instead of a Management Endpoint.
- Added a recommendation on when to Pause after receiving a Pause Control Primitive in the Transmit state.
- Clarified that the behavior of Response Replay Offset is mandatory.

Revision History

Revision Date	Author	Change Description
2019-06-10	Austin Bolen	<ul style="list-style-type: none"> Initial draft. Add missing values from Element Type field in Element Metadata Structure. Clarification to use proper field name in VPD Write section. Grammar fix in Section 1.4 (NVM Subsystem Architectural Model). Require specifying Controller ID in Read NVMe-MI Data Structure reading a Data Structure Type (DTYPE) of 04h (Optional Commands Supported) since the optional commands can vary per controller.
2019-06-10a	Austin Bolen	<ul style="list-style-type: none"> Changed Element Type to only be 5 bits (bits 4:0) and made bit 5 Reserved. Started adding clarifications in Section 5.7 (Read NVMe-MI Data Structure) that the Optionally Supported Command list is for admin commands available out-of-band.
2019-06-11	Austin Bolen	<ul style="list-style-type: none"> Modified grammar in section 1.4 based on feedback. Removed clarifications in Section 5.7 (Read NVMe-MI Data Structure) added in the previous version of this errata document that the Optionally Supported Command list is for admin commands available out-of-band.
2019-07-29	Austin Bolen	<ul style="list-style-type: none"> Fixed errata related to Report All bit where CSTS was missing from the list of Controller selection fields in two places.
2019-10-28	Austin Bolen	<ul style="list-style-type: none"> Removed PCIe link down as something that causes a Management Endpoint reset in 9.3.3 since that is a type of Conventional Reset which is already listed. Removed duplicate sentence in Figure 26. Added clarification to 5.1.3 that changing MTU with outstanding commands results in undefined behavior. Clarified how MsgTag should be set for Replay in 4.2.1.5. Started adding clarification on when to pause after receiving a Pause Control Primitive in 4.2.1.1. Clarified that sending commands to Management Endpoint with one or more paused Command Slots results in undefined behavior in 4.2.1.1.
2019-11-18	Austin Bolen	<ul style="list-style-type: none"> Reworded erratum that changing MTU with outstanding commands results in undefined behavior in 5.1.3. Added clarification that changing MTU after sending a command but before a Replay results in undefined behavior in 5.1.3. Reworded erratum on how MsgTag should be set for Replay in 4.2.1.5. Reworded erratum on when to pause after receiving a Pause Control Primitive in 4.2.1.1. Clarified that it is mandatory to honor Response Replay Offset.
2019-12-02	Austin Bolen	<ul style="list-style-type: none"> Reworded erratum on when to pause after receiving a Pause Control Primitive in 4.2.1.1 again. Clarified that a Command Slot is paused instead of a Management Endpoint.
2020-01-13	Austin Bolen	<ul style="list-style-type: none"> Finished the errata overview section.
2020-01-14	Austin Bolen	<ul style="list-style-type: none"> Editorial updates from workgroup meeting. Accepted changes and deleted comments.
2020-10-16	Austin Bolen	<ul style="list-style-type: none"> Fixed markup colors in Figure 74.
2020-04-09	Jarryd Allison	<ul style="list-style-type: none"> Integrated into NVMe-MI Base Specification version 1.1a
2020-04-13	Mike Allison	<ul style="list-style-type: none"> Made new text blue in color.
2020-04-20	Admin	<ul style="list-style-type: none"> Ratified

Incompatible Changes

- There is a change to require the Controller ID field (CTRLD) in Read NVMe-MI Data Structure when reading a Data Structure Type (DTYPE) of 04h (Optional Commands Supported) since the optional commands can vary per controller.

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

Description of Specification Changes

Modify Figure 140 (Metadata Element Descriptor) as follows:

07:056	Reserved								
054:00	Element Type (ET): This field specifies the type of metadata stored in the descriptor. <table><tr><th>Value</th><th>Definition</th></tr><tr><td>00h</td><td>Reserved</td></tr><tr><td>01h to 017h</td><td>Element types defined by this specification. Controller Metadata Element types are defined in Figure 141. Namespace Metadata Element types are defined in Figure 142.</td></tr><tr><td>18h to 1Fh</td><td>Vendor Specific</td></tr></table>	Value	Definition	00h	Reserved	01h to 017h	Element types defined by this specification. Controller Metadata Element types are defined in Figure 141. Namespace Metadata Element types are defined in Figure 142.	18h to 1Fh	Vendor Specific
Value	Definition								
00h	Reserved								
01h to 017h	Element types defined by this specification. Controller Metadata Element types are defined in Figure 141. Namespace Metadata Element types are defined in Figure 142.								
18h to 1Fh	Vendor Specific								

Modify a portion of 5.12 (VPD Write) as follows:

A VPD Write command with Data Length 0h and no data is valid. The Responder responds with a Success Response.

Modify a portion of 1.4 (NVM Subsystem Architectural Model) as follows:

When using the in-band tunneling mechanism, the architecture and behavior of an NVM Subsystem ~~when the in-band tunneling mechanism is used~~ as defined by the NVM Express specification with extensions defined by this specification.

Modify Figure 88 (Read NVMe-MI Data Structure – NVMe Management Dword 0) as follows:

Bit	Description																
31:24	<p>Data Structure Type (DTYP): This field specifies the data structure to return.</p> <table><tr><th>Value</th><th>Definition</th></tr><tr><td>00h</td><td>NVM Subsystem Information</td></tr><tr><td>01h</td><td>Port Information</td></tr><tr><td>02h</td><td>Controller List</td></tr><tr><td>03h</td><td>Controller Information</td></tr><tr><td>04h</td><td>Optionally Supported Command List</td></tr><tr><td>05h</td><td>Management Endpoint Buffer Command Support List</td></tr><tr><td>06h to FFh</td><td>Reserved</td></tr></table>	Value	Definition	00h	NVM Subsystem Information	01h	Port Information	02h	Controller List	03h	Controller Information	04h	Optionally Supported Command List	05h	Management Endpoint Buffer Command Support List	06h to FFh	Reserved
Value	Definition																
00h	NVM Subsystem Information																
01h	Port Information																
02h	Controller List																
03h	Controller Information																
04h	Optionally Supported Command List																
05h	Management Endpoint Buffer Command Support List																
06h to FFh	Reserved																
23:16	<p>Port Identifier (PORTID): This field contains the identifier of the port whose data structure is returned.</p> <p>If the DTYP field value corresponds to Port Information, then this field contains the Port Identifier whose information is requested.</p> <p>If the DTYP field value corresponds to Management Endpoint Buffer Command Support List, then this field contains the Port Identifier whose information is requested.</p> <p>For all other values of the DTYP field, this field is reserved.</p>																
15:00	<p>Controller Identifier (CTRLID): This field contains the Controller identifier whose data structure is returned.</p> <p>If the DTYP field value corresponds to is 02h (Controller List), or 03h (Controller Information), or 04h (Optionally Supported Command List), then this field contains the Controller Identifier in the NVM Subsystem whose information is requested.</p> <p>If the DTYP field value is 04h (Optionally Supported Command List), then this field is only applicable for commands in the Optionally Supported Command List Data Structure with NMIMT set to a value of 02h (NVMe Admin Command) and shall be ignored for commands with NMIMT set to any value other than 02h.</p> <p>For all other values of the DTYP field, this field is reserved.</p>																

Modify Figure 74 (Controller Health Status Poll – NVMe Management Dword 0) as follows:

Description
<p>Report All (ALL): When this bit is set to '1', a Controller Health Data Structure is returned regardless of the status of the Controller Health Status Changed Flags. The Controller selection fields (SCTLID, MAXRENT, INCF, INCPF, and INCVF) still apply even when this bit is set to '1' but the error selection bits (CWARN, SPARE, PDLU, and CTEMP, and CSTS in Figure 75) do not apply.</p> <p>When this bit is cleared to '0', a Controller Health Data Structure is returned based on the Controller selection fields (SCTLID, MAXRENT, INCF, INCPF, and INCVF) error selection fields (CWARN, SPARE, PDLU, and CTEMP, and CSTS in Figure 75).</p>

Modify a portion of 9.3.3 (Management Endpoint Reset) as follows:

In addition to these conditions, a Management Endpoint associated with a PCI Express port is reset when the PCI Express port is in a **PCI Express conventional reset state**. ~~any of the following states:~~

- ~~• A PCI Express conventional reset; or~~
- ~~• When the PCI Express link is down (i.e., not in the DL_Active state).~~

Modify a portion of Figure 26 (Response Message Status Values) as follows:

Value	Description	Error Response Format Section
...		
04h	<p>Invalid Parameter: Invalid parameter field value. Request Messages received with reserved values in defined fields shall be completed with an Invalid Parameter Error Response. Request Messages received with reserved or unimplemented values in defined fields shall be completed with an Invalid Parameter Error Response. Other error conditions that result in Invalid Parameter Error Response are noted elsewhere in this specification.</p>	4.1.2.2

Modify a portion of 5.2.3 (MCTP Transmission Unit Size (Configuration Identifier 03h)) as follows:

After successful completion of this command, the MCTP Transmission Unit Size for MCTP packets on the specified port is updated to the specified size for future Command Messages. A Management Controller should not change this configuration while there are other commands outstanding. **Changing this configuration while there are other Request Messages outstanding results in undefined behavior. If a Request Message is sent with a given MCTP Transmission Unit Size, then issuing a Replay Control Primitive after changing the MCTP Transmission Unit Size to a different value results in undefined behavior.**

Modify a portion of 4.2.1.1 (Pause) as follows:

Transmit: The Pause Control Primitive sets the Pause Flag to '1' (refer to section 4.2.1.4) suspending transmission of ~~MCTP Response Messages~~ packets associated on a packet boundary ~~with the Command in the Command Slot~~. The Management Endpoint should pause transmission as soon as possible after receiving a Pause Control Primitive.

The Management Endpoint shall transmit a Response Message with success status after receiving the Pause primitive. It is not an error to issue a Pause Control Primitive when a Command Slot is already paused.

While the Pause Flag is set to '1', the Management Endpoint disables the timeout waiting for packet timer and does not transmit responses to commands. The timeout waiting for a packet is the lesser of 100 ms or the time defined in the appropriate MCTP transport binding specification. The Management Controller should not send ~~Commands~~ Messages to a Command Slot that ~~while a Management Endpoint~~ is paused.

Modify a portion of Figure 40 (Replay Control Primitive Request Fields) as follows:

Byte	Description						
07:06	Control Primitive Specific Parameter (CPSP): This field is used to pass Control Primitive specific parameter information.						
	<table><tr><th>Bits</th><th>Description</th></tr><tr><td>15:08</td><td>Reserved</td></tr><tr><td>07:00</td><td>Response Replay Offset (RRO): This field specifies the starting packet number from which the Response Message associated with the last Command Message processed in the Command Slot shall should be replayed. This is a 0's based value. When this field is cleared to 0h, the first packet of the associated Response Message is the first packet replayed. If this field specifies an offset that is beyond the length of the Response Message, then processing of the Control Primitive is aborted and the Management Endpoint transmits an Invalid Parameter Error Response.</td></tr></table>	Bits	Description	15:08	Reserved	07:00	Response Replay Offset (RRO): This field specifies the starting packet number from which the Response Message associated with the last Command Message processed in the Command Slot shall should be replayed. This is a 0's based value. When this field is cleared to 0h, the first packet of the associated Response Message is the first packet replayed. If this field specifies an offset that is beyond the length of the Response Message, then processing of the Control Primitive is aborted and the Management Endpoint transmits an Invalid Parameter Error Response.
	Bits	Description					
15:08	Reserved						
07:00	Response Replay Offset (RRO): This field specifies the starting packet number from which the Response Message associated with the last Command Message processed in the Command Slot shall should be replayed. This is a 0's based value. When this field is cleared to 0h, the first packet of the associated Response Message is the first packet replayed. If this field specifies an offset that is beyond the length of the Response Message, then processing of the Control Primitive is aborted and the Management Endpoint transmits an Invalid Parameter Error Response.						