



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	6029 – Allow Parameter Error Location not to be ordered
Change Date	2021.02.22
Builds on Specification	NVM Express Management Interface 1.1b

Technical Proposal Author(s)

Name	Company
Michael Allison	Intel
Myron Loewen	Intel

This proposal is to remove the ordering rules for reporting Parameter Error Location and allow it to be vendor specific which error is reported when multiple errors exist.

Revision History

Revision Date	Change Description
11/04/2020	Initial Draft
11/09/2020	Minor editorial changes per Management Interface Task Group review.
12/17/2020	Accepted all changes and removed all comments for member review.
1/27/2021	Removed the double negative in section 4.2.3.
2/14/2021	Integrated into the NVMe-Management Interface Specification, Revision 1.1.
2/22/2021	Removed all comments, accepted all changes, and converted references/cross-references to text.

Description for NVMe-Mi Changes Document

- Feature Enhancements
 - Invalid Parameter Error Response enhancement
 - Removed the ordering restriction on if multiple parameter errors exists in a Request Message. (Optional)
 - Updated parameter error text to be explicit on the field to be reported in error.

Description of Specification Changes

Markup Conventions:

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

Modify a portion of section 3.1.1 as follows:

3.1.1 Message Fields

..

The Management Endpoint Buffer (MEB) bit in the Message Header specifies whether Message Data is contained in the associated Message Data field of an NVMe-MI Message or in the Management Endpoint Buffer. This bit should only be set to '1' in Command Messages that support Management Endpoint Buffer operation (i.e., those listed in the Management Endpoint Buffer Supported Command List data structure). If the MEB bit is set to '1' it is an error to set this bit in any other Command Message, then the Management Endpoint shall respond ~~and when this occurs it causes the Command Message to complete~~ with an Invalid Parameter Error Response with the PEL field indicating the MEB bit.

Modify a portion of section 4.1.2.2 and Figure 29 as follows:

4.1.2.2 Invalid Parameter Error Response

An Invalid Parameter Error Response is generated for Error Responses where the Status field is set to Invalid Parameter. The format of an Invalid Parameter Error Response is shown in Figure 28 and the response specific fields are summarized in Figure 29.

Unless otherwise specified, if multiple invalid parameter errors exist in a Request Message, then the Management Endpoint selects the invalid parameter that is returned in the Invalid Parameter Error Response.

Figure 28: Invalid Parameter Error Response

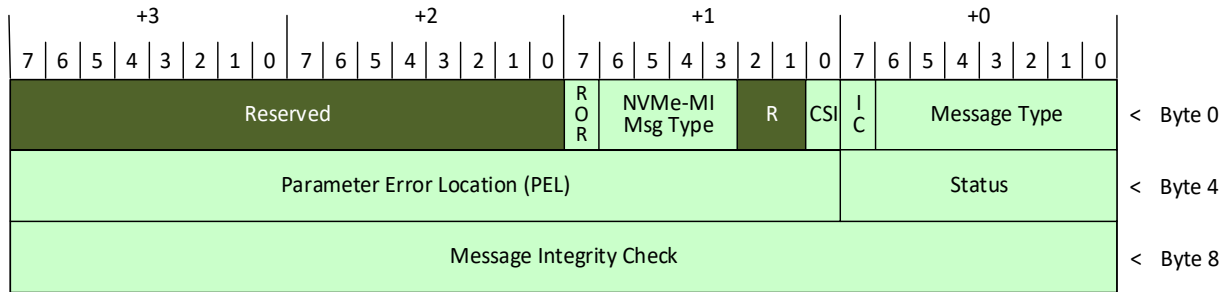


Figure 29: Invalid Parameter Error Response Fields

Byte	Description								
7:5	<p>Parameter Error Location (PEL): This field indicates the byte and bit of the request parameter within the Request Message that contains the an first invalid parameter -(i.e., the invalid parameter with the lowest byte and bit).</p> <p>If the invalid parameter spans multiple bytes or bits, then the location indicates the first byte and bit of the parameter.</p> <table> <tr> <th>Bits</th><th>Description</th></tr> <tr> <td>23:08</td><td>Byte in the Request Message of the parameter that contained the error. If the error is beyond byte 65,535, then the value 65,535 is reported in this field.</td></tr> <tr> <td>07:03</td><td>Reserved</td></tr> <tr> <td>02:00</td><td>Bit in the Request Message of the parameter that contained the error. Valid values are 0 to 7.</td></tr> </table>	Bits	Description	23:08	Byte in the Request Message of the parameter that contained the error. If the error is beyond byte 65,535, then the value 65,535 is reported in this field.	07:03	Reserved	02:00	Bit in the Request Message of the parameter that contained the error. Valid values are 0 to 7.
Bits	Description								
23:08	Byte in the Request Message of the parameter that contained the error. If the error is beyond byte 65,535, then the value 65,535 is reported in this field.								
07:03	Reserved								
02:00	Bit in the Request Message of the parameter that contained the error. Valid values are 0 to 7.								

Modify a portion of section 4.2.1.1 as follows:

4.2.1.1 Pause

The Pause Control Primitive is used to suspend response transmission and suspend the timeout waiting for packet for both Command Slots in a Management Endpoint. The CSI bit in a Pause Control Primitive is not used and shall be cleared to 0h. If the CSI bit is set to '1', then the Management Endpoint should transmit an Invalid Parameter Error Response **with the PEL field indicating the CSI bit.**

Modify a portion of figure 40 in section 4.2.1.5 as follows:

4.2.1.5 Replay

...

Figure 40: Replay Control Primitive Request Fields

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><p>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 be replayed.</p><p>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.</p><p>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 shall transmit an Invalid Parameter Error Response with the PEL field indicating this field.</p></td></tr></table>	Bits	Description	15:08	Reserved	07:00	<p>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 be replayed.</p> <p>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.</p> <p>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 shall transmit an Invalid Parameter Error Response with the PEL field indicating this field.</p>
	Bits	Description					
15:08	Reserved						
07:00	<p>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 be replayed.</p> <p>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.</p> <p>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 shall transmit an Invalid Parameter Error Response with the PEL field indicating this field.</p>						

Modify a portion of section 4.2.3 as follows:

4.2.3 Management Endpoint Buffer

...

If the Management Endpoint Buffer (MEB) bit is set to '1' in a Command Message that normally contains Request Data, then ~~no Request Data is~~ shall not be transferred in the Command Message itself and the required Request Data ~~is instead~~ shall be transferred from the Management Endpoint Buffer. The Request Data starts at a zero offset from the start of the Management Endpoint Buffer. ~~If the MEB bit is set to '1' in a Command Message that normally contains Request Data, then the Command Message shall contain no Request Data.~~ If ~~the~~ that Command Message contains Request Data or ~~is one that~~ does not support Request Data, then the Management Endpoint responds with an Invalid Parameter Error Response with the PEL field indicating the ~~the~~ Request Data field of the Command Message.

Modify a portion of section 5.2 as follows:

5.2 Configuration Set

...

NVMe-MI Configuration Identifiers are listed in Figure 62. Specifying a reserved identifier in the Configuration Identifier field causes the command to complete with an Invalid Parameter Error Response with the PEL field indicating the Configuration Identifier field.

Modify a portion of section 5.2.1 as follows:

5.2.1 SMBus/I2C Frequency (Configuration Identifier 01h)

...

If the specified frequency is not supported, then the Management Endpoint shall respond with an Invalid Parameter Error Response with the PEL field indicating the SMBus/I2C Frequency field. ~~or~~ If the Port Identifier specified is not an SMBus/I2C port, then the Management Endpoint shall respond with an Invalid Parameter Error Response with the PEL field indicating the Port Identifier field.

Modify a portion of section 5.2.3 as follows:

5.2.3 MCTP Transmission Unit Size (Configuration Identifier 03h)

...

If the specified MCTP Transmission Unit Size is not supported, then the Management Endpoint shall abort the command and send a Response Message with an Invalid Parameter Error Response with the PEL field indicating the MCTP Transmission Unit Size field. ~~or~~ If the Port Identifier specified is not valid, then the Management Endpoint shall abort the command and send a Response Message with an Invalid Parameter Error Response with the PEL field indicating the Port Identifier field.

Modify a portion of section 5.4 and Figure 82 as follows:

5.4 Management Endpoint Buffer Read

...

If the Data Offset (DOFST) field is greater than or equal to the size of the Management Endpoint Buffer, then the Management Endpoint responds with an Invalid Parameter Error Response with the PEL field indicating ~~.The parameter with the error in this case is~~ the DOFST field. If the DOFST field is less than the size of the Management Endpoint Buffer and the sum of the DOFST and DLEN fields is greater than or equal to size of the Management Endpoint Buffer, then the Management Endpoint responds with an Invalid Parameter Error Response with the PEL field indicating ~~.The parameter with the error in this case is~~ the DLEN field.

...

Figure 82: Management Endpoint Buffer Read – NVMe Management Dword 1

Bit	Description
31:16	Reserved
15:00	<p>Data Length (DLEN): This field specifies the length, in bytes, to be transferred from the Management Endpoint Buffer starting at the byte offset specified by DOFST and returned in the Response Data. Specifying a DLEN field value that is greater than the maximum supported Response Data size results in an Invalid Parameter Error Response with the PEL field indicating this field.</p> <p>A Data Length value of 0h and no data is valid. The Management Endpoint responds with a Success Response and no Response Data.</p>

Modify a portion of section 5.5 and Figure 85 as follows:

5.5 Management Endpoint Buffer Write

...

If the Data Offset (DOFST) field is greater than or equal to the size of the Management Endpoint Buffer, then the Management Endpoint responds with an Invalid Parameter Error Response with the PEL field indicating ~~.The parameter with the error in this case is~~ the DOFST field. If the DOFST field is less than the size of the Management Endpoint Buffer and the sum of the DOFST and DLEN fields is greater than or equal to size of the Management Endpoint Buffer, then the Management Endpoint responds with an Invalid Parameter Error Response with the PEL field indicating ~~.The parameter with the error in this case is~~ the DLEN field.

...

Figure 85: Management Endpoint Buffer Write – NVMe Management Dword 1

Bit	Description
31:16	Reserved
15:00	<p>Data Length (DLEN): This field specifies the length, in bytes, to be transferred from the Request Data to the Management Endpoint Buffer starting at the byte offset specified by DOFST. Specifying a DLEN field value that is greater than the maximum supported Response Data size results in an Invalid Parameter Error Response with the PEL field indicating this field.</p> <p>A Data Length value of 0h specifies that no data shall be transferred. This condition shall not be considered an error.</p>

Modify a portion of section 5.9 as follows:

5.9 SES Receive

...

The Page Code (PCODE) field specifies the SES status type diagnostic page to be retrieved. Refer to SES-3 for a list and description of SES diagnostic pages. If the PCODE field specifies a reserved value, an unsupported value, or a value that only corresponds to an SES control type diagnostic page, then the Responder responds with an Invalid Parameter Error Response [with the PEL field indicating the PCODE field](#).

The Allocation Length (ALENGTH) field specifies the maximum length of the Response Data field in the Response Message and is used to limit the maximum amount of SES diagnostic page data that may be returned. The length of the Response Data field shall be the total length of the SES diagnostic page specified by the PCODE field or the number of bytes specified by the ALENGTH field (i.e., the SES diagnostic page is truncated), whichever is less. When the SES diagnostic page is truncated, the value of fields within the SES diagnostic page are not altered to reflect the truncation.

All errors are detected and reported while servicing the SES Receive command and reported via an Error Response. If an invalid field is detected in an SES Receive command, then the Responder responds with an Invalid Parameter Error Response [with the PEL field indicating the invalid field](#). If a condition occurs that in SES-3 results in a CHECK CONDITION, then the Responder responds with an Error Response. The mapping of Error Response Status values to SES-3 sense keys and additional sense codes is shown in Figure 12.

Modify a portion of section 5.10 as follows:

5.10 SES Send

...

Unlike the SES Receive command that specifies the page code of the SES status diagnostic page being retrieved, the SES Send command specifies the page code of the SES control type diagnostic page that is being transferred in the SES control type diagnostic page itself. Refer to SES-3 for a list and description of SES control type diagnostic pages. If the [pPage eCode \(PCODE\) field](#) in the SES control type diagnostic page specifies a reserved value, an unsupported value, or a value that only corresponds to an SES status diagnostic page, then the Responder responds with an Invalid Parameter Error Response [with the PEL field indicating the PCODE field](#).

The SES Send command does not use NVMe Management Dword 0 or the NVMe Management Response field. All of these are reserved.

All errors are detected and reported while processing the SES Send command and reported via an Error Response. If an invalid field is detected in the SES control type diagnostic page data transferred by an SES Send command, then the Responder responds with an Invalid Parameter Error Response [with the PEL](#)

field indicating the invalid field. If a condition occurs that in SES-3 results in a CHECK CONDITION, then the Responder responds with an Error Response. The mapping of Response Message Status values to SES-3 sense keys and additional sense codes is shown in Figure 12.

Modify a portion of section 5.11 as follows:

5.11 VPD Read

...

A VPD Read command with length 0 and no data is valid. The Responder responds with a Success Response and no Response Data. If the Data Length plus Data Offset fields are greater than the size of the VPD, then the Responder does not return the VPD contents and responds with an Invalid Parameter Error Response with the PEL field indicating the Data Length field.

Modify a portion of section 5.12 as follows:

5.12 VPD Write

...

If the Data Length plus Data Offset fields are greater than the size of the VPD, then the Responder does not write to the VPD and responds with an Invalid Parameter Error Response with the PEL field indicating the Data Length field.

Modify a portion of section 6 as follows:

6 NVM Express Admin Command Set

The NVM Express Admin Command Set allows NVMe Admin Commands to be issued to any Controller in the NVM Subsystem using the out-of-band mechanism. Figure 110 shows NVM Express Admin Commands that are mandatory, optional, and prohibited for an NVMe Storage Device and an NVMe Enclosure using the out-of-band mechanism. All NVM Express Admin Commands are prohibited using the in-band tunneling mechanism. The commands are defined in the NVM Express specification. If an NVMe Admin Command is issued in a Request Message that is a prohibited command in Figure 110, the Management Endpoint shall return an Invalid Parameter Error Response with ~~Parameter-Error-Location-pointing-field~~ indicating to the NVMe opcode. Future revisions of this specification may add additional commands to Figure 110. The NVM Express Admin Command Set is only applicable in the out-of-band mechanism and is prohibited in the in-band tunneling mechanism.

...

Figure 112: NVMe Admin Command Request Description

Byte	Description
...	

Figure 112: NVMe Admin Command Request Description

Byte	Description
31:28	<p>Data Offset (DOFST): For commands that transmit data from the Management Controller to the Management Endpoint (i.e., the Request Data field in the Request Message has non-zero length) or do not transmit data, this field shall be cleared to 0h. If this field is not 0h, then the Management Endpoint shall return an Invalid Parameter Error Response with the PEL field indicating this field.</p> <p>For commands that transmit data from the Management Endpoint to the Management Controller (i.e., the Response Data field in the Response Message has non-zero length), this field specifies the starting offset, in bytes, of the portion of data contained in the NVMe Admin Command completion data that is returned starting at byte offset 0h of the Response Data field in the Response Message.</p> <p>Bits 0 and 1 of this field shall be cleared to '0'.</p>
35:32	<p>Data Length (DLEN): For commands that do not transmit data in either the Request Message or Response Message, this field shall be cleared to 0h. If this field is not 0h, then the Management Endpoint shall return an Invalid Parameter Error Response with the PEL field indicating this field.</p> <p>For commands that transmit data from the Management Controller to the Management Endpoint (i.e., the Request Data field in the Request Message has non-zero length), this field specifies the length, in bytes, of the data contained in the Request Data field in the Request Message.</p> <p>For commands that transmit data from the Management Endpoint to the Management Controller (i.e., the Response Data field in the Response Message has non-zero length), this field specifies the length, in bytes, of the portion of data contained in the NVMe Admin Command completion data that is returned in the Response Data field in the Response Message.</p> <p>Bits 0 and 1 of this field shall be cleared to '0'. This field shall be less than or equal to 4,096.</p>

Modify a portion of section 6.1 as follows:

6.1 Request and Response Data

NVMe Admin Commands may contain data as part of the Command Message. This data is passed in the Request Data field instead of using PRP Lists or SGL segments. The PRP Entry 2 (PRP2) and Metadata Pointer (MPTR) fields within the NVMe Admin Commands are reserved.

If there is no data sent with the NVMe Admin Command (e.g., the Data Transfer subfield for the opcode is 00b), then the Data Offset and Data Length fields shall be cleared to 0h.

If there is data sent with the NVMe Admin Command (i.e., the Data Transfer subfield for the opcode is 01b), then the Data Offset field shall be 0h and the Data Length field shall be set to the length of the Request Data required by the command. If the Data Length field does not correspond to the required length, the Management Endpoint shall respond with an Invalid Parameter Error Response [with the PEL field indicating the Data Length field](#).

If there is Response Data expected in the Response Message in the completion of the NVMe Admin Command (i.e., the Data Transfer subfield in the corresponding NVMe Admin Command for the opcode is 10b), then the Data Offset and Data Length fields describe the portion of the NVMe Admin Command completion data that is transferred in the Response Message. Any remaining data not transferred in the Response Message is discarded by the Management Endpoint as shown in Figure 115. If the Data Length plus Data Offset fields are greater than the size of the NVMe Admin Command completion data, the Management Endpoint should respond with an Invalid Parameter Error Response [with the PEL field indicating the Data Offset field](#).

Modify a portion of section 7.1 as follows:

7.1 PCIe Configuration Read

...

If the sum of the Offset and Length fields fall outside of PCI configuration space, then the Management Endpoint responds with an Invalid Parameter Error Response [with the PEL field indicating](#). ~~The parameter with the error in this case is always~~ the Offset field.

Modify a portion of section 7.2 as follows:

7.2 PCIe Configuration Write

..

If the sum of the Offset and Length fields fall outside of PCI configuration space, then the Management Endpoint responds with an Invalid Parameter Error Response [with the PEL field indicating](#). ~~The parameter with the error in this case is always~~ the Offset field.

Modify a portion of section 7.3 as follows:

7.3 PCIe I/O Read

If the Base Address Register field does not correspond to an I/O BAR implemented by the specified NVMe Controller, then the Management Endpoint responds with an Invalid Parameter Error Response [with the PEL field indicating the Base Address Register field](#).

If the sum of the Offset and Length fields fall outside the address range of the BAR specified by the Base Address Register field, then the Management Endpoint responds with an Invalid Parameter Error Response [with the PEL field indicating](#). ~~The parameter with the error in this case is always~~ the Offset field.

Modify a portion of section 7.4 as follows:

7.4 PCIe I/O Write

...

When this command is completed successfully, PCI I/O space associated with the NVMe Controller specified by Controller ID is written with the data contained in the Request Data field. The Offset field specifies the starting write offset in PCIe I/O address space specified by the Base Address Register field. The Length field specifies the number of bytes to be written. The Request Data field is always an integral number of dwords and is equal to the Length field rounded up to the next dword. If Length is not an integral number of dwords, then unused padding bytes are discarded.

If the Base Address Register field does not correspond to an I/O BAR implemented by the specified NVMe Controller, then the Management Endpoint responds with an Invalid Parameter Error Response [with the PEL field indicating the Base Address Register field](#).

If the sum of the Offset and Length fields fall outside the address range of the BAR specified by the Base Address Register field, then the Management Endpoint responds with an Invalid Parameter Error Response [with the PEL field indicating](#). ~~The parameter with the error in this case is always~~ the Offset field.

Modify a portion of section 7.5 as follows:

7.5 PCIe Memory Read

...

If the Base Address Register field does not correspond to one implemented by the specified NVMe Controller, or the address range specified by the Base Address Range is not a memory region, then the Management Endpoint responds with an Invalid Parameter Error Response [with the PEL field indicating the Base Address Register field](#).

If the sum of the Offset and Length fields fall outside the address range specified by the Base Address Register field, then the Management Endpoint responds with an Invalid Parameter Error Response [with the PEL field indicating](#).~~The parameter with the error in this case is always~~ the Offset field.

Modify a portion of section 7.6 as follows:

7.6 PCIe Memory Write

...

If the Base Address Register field does not correspond to one implemented by the specified NVMe Controller, or the address range specified by the Base Address Range is not a memory region, then the Management Endpoint responds with an Invalid Parameter Error Response [with the PEL field indicating the Base Address Register field](#).

If the sum of the Offset and Length fields fall outside the address range of the BAR specified by the Base Address Register field, then the Management Endpoint responds with an Invalid Parameter Error Response [with the PEL field indicating](#).~~The parameter with the error in this case is always~~ the Offset field.