



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
info@nvmexpress.org

NVM Express Technical Proposal for New Feature

Technical Proposal ID	6010b – Command Initiated Auto Pause
Change Date	02/22/2021
Builds on Specification	NVM Express Management Interface 1.1

Technical Proposal Author(s)

Name	Company
Kimito Sakata	Oracle
Mike Allison	Intel
John Geldman	KioXia

Feature to improve efficiency and performance for systems based on MUX SMBus architecture. The feature defines a bit in the command header that will tell the device endpoint to go into a PAUSE primitive state as if the host endpoint sent the PAUSE primitive after the command header. The device endpoint will behave as if the PAUSE primitive has been sent and aside from the lack of PAUSE primitive transaction, will behave as if the PAUSE primitive was sent. This would free the host endpoint to go switch the MUX channel and service other devices if the command overhead is big enough so that service can be done to other devices. The host after switching back to the MUX channel will continue by sending the RESUME primitive. Aside from the PAUSE primitive being sent to the device endpoint and the lack of PAUSE response data returned by the device endpoint, everything else should behave as if the PAUSE primitive was sent by the host endpoint after the command was sent to the device endpoint.

Revision History

Revision Date	Change Description
02/27/2019	<ul style="list-style-type: none">Initial draft.
06/13/2019	<ul style="list-style-type: none">Integration
06/17/2019	<ul style="list-style-type: none">Moved reference from 4.2.1 to 4.2.1.1.
04/28/2020	<ul style="list-style-type: none">Added text to update the NVMe-MI messages diagrams that allow the CIAP bit to show the bit.
06/06/2020	<ul style="list-style-type: none">Moved the text defining the CIAP bit so the it aligned with the other bit definitions in relation to the order in the diagram.
6/15/2020	<ul style="list-style-type: none">Aligned to the same text for the MEB bit in TP 6020. Added more figures to the list of figures to add the CIAP bit.
7/13/2020	<ul style="list-style-type: none">Ready for member review
8/17/2020	<ul style="list-style-type: none">Clarified the parameter in error when a Invalid Parameter Error Response occurs.
9/27/2020	<ul style="list-style-type: none">Integrated into the NVMe Management Interface Specification, Revision 1.1.
11/4/2020	<ul style="list-style-type: none">Incorporated new phases for Invalid Parameter Error Responses.
12/17/2020	<ul style="list-style-type: none">Accepted all changes and removed all comments for member review.
02/14/2021	<ul style="list-style-type: none">Integrated into the NVMe-Management Interface Specification, Revision 1.1.
02/22/2021	<ul style="list-style-type: none">Removed all comments, accepted all changes, and converted references/cross-references to text.

Description of Specification Changes

Section 3.1.1 Message Fields, add the paragraph in red below:

The Request or Response (ROR) bit in the Message Header specifies whether the NVMe-MI Message is associated with a Request Message or a Response Message. The NVMe-MI Message Type (NMIMT) field specifies whether the Request Message is a Control Primitive or a specific type of Command Message (refer to Figure 23). Finally, the Command Slot Identifier (CSI) bit specifies the Command Slot with which the NVMe-MI Message is associated in the out-of-band mechanism. Refer to section 4.2 for additional information about Command Slots.

~~The Command Initiated Auto Pause (CIAP) bit in the Message Header of a Command Message specifies whether or not the Management Endpoint is automatically paused when a Command Message enters the Process state. A Command Message with the CIAP bit set to '1' shall be treated by the Management Endpoint as if an implicit Pause Control Primitive, as described in section 4.2.1.1, was received in the Process state with the exception that the Management Endpoint shall not transmit a Control Primitive Response Message. The Command Initiated Auto Pause Supported (CIAPS) bit in Figure 91 indicates if the port supports the Command Initiated Auto Pause (CIAP) bit in Command Messages.~~

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 in Command Messages that support Management Endpoint Buffer operation (i.e., those listed in the Management Endpoint Buffer Supported Command List data structure). It is an error to set this bit in any other Command Message and when this occurs it causes the Command Message to complete with an Invalid Parameter Error Response.

~~The Command Initiated Auto Pause (CIAP) bit in the Message Header of a Command Message specifies whether or not the Management Endpoint is automatically paused when a Command Message enters the~~

Process state. A Command Message with the CIAP bit set to '1' shall be treated by the Management Endpoint as if an implicit Pause Control Primitive, as described in section 4.2.1.1, was received in the Process state with the exception that the Management Endpoint shall not transmit a Control Primitive Response Message. The Command Initiated Auto Pause Supported (CIAPS) bit in Figure 91 indicates if the port supports the Command Initiated Auto Pause (CIAP) bit in Command Messages.

Modify Figure 18 as shown in red:

Figure 18: NVMe-MI Message Fields

1

NVMe-MI Message Parameters (NMP): This field contains parameters applicable to the NVMe-MI Message.		
Bits	Description	
7	Request or Response (ROR): This bit indicates whether the message is a Request Message or Response Message. This bit is cleared to '0' for Request Messages. This bit is set to '1' for Response Messages.	
6:3	NVMe-MI Message Type (NMIMT): This field specifies the NVMe-MI Message Type. Refer to the sections referenced in the table below for details about each NVMe-MI Message Type and whether they apply to the out-of-band mechanism, the in-band tunneling mechanism, or both.	
	Value	Description
	0h	Control Primitive
	1h	NVMe-MI Command
	2h	NVMe Admin Command
	3h	Reserved
	4h	PCIe Command
5h to Fh	Reserved	
2:1	Reserved	
0	Command Slot Identifier (CSI): This bit indicates the Command Slot with which the NVMe-MI Message is associated. For Request Messages this bit indicates the Command Slot with which the Request Message is associated. For Response Messages, this bit indicates the Command Slot associated with the Request Message with which the Response Message is associated. This bit is only applicable to NVMe-MI Messages in the out-of-band mechanism. This bit is reserved for NVMe-MI Messages in the in-band tunneling mechanism.	
	Value	Description
	0b	Command Slot 0
	1b	Command Slot 1

9

2

Bits	Description						
7:2	Reserved						
1	<p>Command Initiated Auto Pause (CIAP): If this bit is set to '1' in a Command Message, the Management Endpoint shall be automatically paused when the Command Message enters the Process state. If this bit is cleared to '0' in a Command Message, the Management Endpoint shall not be automatically paused when the Command Message enters the Process state.</p> <p>This bit is only valid for Command Messages sent using the out-of-band mechanism and is reserved for all other types of NVMe-MI Messages.</p> <p>If this bit is set to '1' for any type of NVMe-MI Message received by a Management Endpoint other than a Command Message, then an Invalid Parameter Error Response with the PEL field indicating this bit shall be returned.</p> <p>If this bit is set to '1' in the in-band tunneling mechanism, then an Invalid Parameter Error Response with the PEL field indicating this bit shall be returned.</p>						
0	<p>Management Endpoint Buffer (MEB): This bit indicates whether the Message Data is contained in the Message Data field of this NVMe-MI Message or in the Management Endpoint Buffer. Refer to section 3.1.</p> <table><tr><th>Value</th><th>Description</th></tr><tr><td>0b</td><td>The Message Data is contained in the Message Data of this NVMe-MI Message.</td></tr><tr><td>1b</td><td>The Message Data is contained in the Management Endpoint Buffer.</td></tr></table>	Value	Description	0b	The Message Data is contained in the Message Data of this NVMe-MI Message.	1b	The Message Data is contained in the Management Endpoint Buffer.
Value	Description						
0b	The Message Data is contained in the Message Data of this NVMe-MI Message.						
1b	The Message Data is contained in the Management Endpoint Buffer.						

Modify Figure 91 as shown in red:

Figure 91: Port Information Data Structure

Byte	Description										
00	<p>Port Type: Specifies the port type.</p> <table border="1"> <thead> <tr> <th>Value</th><th>Definition</th></tr> </thead> <tbody> <tr> <td>0h</td><td>Inactive</td></tr> <tr> <td>1h</td><td>PCIe</td></tr> <tr> <td>2h</td><td>SMBus</td></tr> <tr> <td>3h to FFh</td><td>Reserved</td></tr> </tbody> </table>	Value	Definition	0h	Inactive	1h	PCIe	2h	SMBus	3h to FFh	Reserved
Value	Definition										
0h	Inactive										
1h	PCIe										
2h	SMBus										
3h to FFh	Reserved										
01	<p>Reserved</p> <p>Port Capabilities: This field contains information about the capabilities of the port.</p> <table border="1"> <thead> <tr> <th>Bit</th><th>Description</th></tr> </thead> <tbody> <tr> <td>7:1</td><td>Reserved</td></tr> <tr> <td>0</td><td>Command Initiated Auto Pause Supported (CIAPS): If this bit is set to '1', then the Command Initiated Auto Pause (CIAP) bit is supported in Command Messages on this port. If this bit is cleared to '0', then the CIAP bit is not supported in Command Messages on this port.</td></tr> </tbody> </table>	Bit	Description	7:1	Reserved	0	Command Initiated Auto Pause Supported (CIAPS): If this bit is set to '1', then the Command Initiated Auto Pause (CIAP) bit is supported in Command Messages on this port. If this bit is cleared to '0', then the CIAP bit is not supported in Command Messages on this port.				
Bit	Description										
7:1	Reserved										
0	Command Initiated Auto Pause Supported (CIAPS): If this bit is set to '1', then the Command Initiated Auto Pause (CIAP) bit is supported in Command Messages on this port. If this bit is cleared to '0', then the CIAP bit is not supported in Command Messages on this port.										
03:02	<p>Maximum MCTP Transmission Unit Size: The maximum MCTP Transmission Unit size the port is capable of sending and receiving.</p> <p>If the port does not support MCTP, then this field shall be cleared to 0h.</p> <p>If the port type is PCIe and the port supports MCTP, then this field shall be set to a value between 64 bytes and the PCIe Max Payload Size supported minus 4, inclusive. All PCIe ports within an NVM Subsystem should report the same value in this field.</p> <p>If the port type is SMBus and the port supports MCTP, then this field shall be set to a value between 64 bytes and 250 bytes, inclusive.</p>										

Figure 91: Port Information Data Structure

Byte	Description
07:04	This field specifies the size of the Management Endpoint Buffer in bytes when a Management Endpoint Buffer is supported. A value of 0h in this field indicates that the Management Endpoint does not support a Management Endpoint Buffer.
31:08	Port Type Specific (refer to Figure 92 and Figure 93)

Modify the following figures to include the CIAP bit in bit 1 of byte 2 of Command Messages (to the left of the MEB bit in byte 2):

- ***Figure 17***
- ***Figure 42***
- ***Figure 48***
- ***Figure 53***
- ***Figure 111***
- ***Figure 117***