



LEGAL NOTICE:

© **Copyright 2007 - 2019 NVM Express, Inc. ALL RIGHTS RESERVED.**

This NVM Express Management Interface revision 1.0a 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 1.0a 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 - 2019 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 NVM Express Administration
3855 SW 153rd Drive
Beaverton, OR 97003
admin@nvmexpress.org

NVM Express Technical Proposal for New Feature

Technical Proposal ID	6005 – NVMe 1.3 Alignment
Change Date	3/20/2019
Builds on Specification	NVM Express Management Interface 1.0a

Technical Proposal Author(s)

Name	Company
Peter Onufryk	Microsemi
Michael Allison	Intel
Austin Bolen	Dell EMC

This proposal is to align the NVMe-MI specification with the NVMe 1.3 specification addressing the following items:

- Identify the new Admin commands allowed over NVMe-MI
- Define commands allowed during Sanitize

Revision History

Revision Date	Change Description
9/4/2017	<ul style="list-style-type: none">Initial draft.
9/19/2017	<ul style="list-style-type: none">Incorporate feedback from Austin Bolen:
9/20/2017	<ul style="list-style-type: none">Corrected the legal notice version numbersIncorporated more feedback from Austin Bolen.
9/21/2017	<ul style="list-style-type: none">Added clarification to the definition of a FRU.Removed unused superscript in Figure XClarified the out-of-band prohibited NVMe Admin Command note.
9/25/2017	<ul style="list-style-type: none">Added the definition of an NVMe storage device.Moved Sanitize operation to section 6.2Removed sending Invalid Parameter if optional NVMe Admin command is not implement.
10/2/2017	<ul style="list-style-type: none">Adjusted the definition of an NVMe Storage Device.
10/9/2017	<ul style="list-style-type: none">Removed definition of NVMe Storage Device since it is defined in the SES TP (TP 001).Refer to NVMe Storage Device definition from the FRU definition.
10/16/2017	<ul style="list-style-type: none">Define NVMe Storage Device FRURedefined FRU
10/23/2017	<ul style="list-style-type: none">Modified definitions of FRU and NVMe Storage Device FRU
10/30/2017	<ul style="list-style-type: none">Split NVMe Storage Device FRU and NVMe Storage Device into two separate definitions (NVMe Storage Device does not necessarily need to be a FRU).
10/6/2017	<ul style="list-style-type: none">Removed extra spaceExisting text that was marked for deletion was not correct. Missing a letter from the original text and an extra space needed to be marked for deletion.
12/17/2017	<ul style="list-style-type: none">Split definition of NVMe Storage Device from that of an NVMe Storage Device Field-Replaceable Unit.Added editor note to change all instances of "NVMe storage device" to "NVMe Storage Device."Added Section 1.4 changes to align with NVMe Storage Device and NVMe Storage Device Field Replaceable Unit definitions.
1/28/2018	<ul style="list-style-type: none">Editorial changes to align with NVMe TP numbering, updated NVMe administration, and updated year.
11/18/2018	<ul style="list-style-type: none">Updated TP to align with TP 6003.Removed FRU, NVMe Storage Device, and NVMe Storage Device FRU definitions from this TP since they are now defined in TP 6003.
3/20/2019	<ul style="list-style-type: none">Ratified

Description of Specification Changes

Editor Note: **Highlighted text** is identifying references that need to be updated or added.

Editor Note: Change all instances of NVMe storage device to NVMe Storage Device.

Editor Note: This TP should be applied prior to the changes in TP 6003. Where changes in this TP and TP 6003 conflict, the changes in TP 6003 take precedence.

Update Section 1.4 as follows:

1.4 Architectural Model

An NVMe ~~s~~Storage ~~d~~Device, such as a PCIe SSD, that implements this specification, consists of an NVM

Subsystem with one or more Management Endpoints. There may be up to one Management Endpoint per PCIe port and SMBus/I2C port. Each Management Endpoint has a Port Identifier that is less than or equal to the Number of Ports (NUMP) field value in the NVM Subsystem Information Data Structure. The Port Identifier for a PCIe port is the same as the Port Number field in the PCIe Link Capabilities Register.

An NVMe Storage Device that is a Field-Replaceable Unit is a physical component, device, or assembly that an end user or technician can remove and replace without having to replace the entire system in which it is contained. Examples of NVMe Storage Device Field-Replaceable Units include a U.2 PCIe SSD, a PCI Express Card Electromechanical add-in card, and an M.2 module. The FRU referenced by the FRU Globally Unique Identifier (FGUID) field in the NVMe Express Specification shall be an NVMe Storage Device Field-Replaceable Unit.

NVMe-MI supports Vital Product Data (VPD) that utilizes the format defined in the IPMI Platform Management FRU Information Storage Definition and is stored in a FRU Information Device. The FRU Information Device may be implemented in the NVM Subsystem, in an external device (e.g., serial EEPROM), or a combination of the two. The VPD is accessible over any port that supports NVMe-MI using MCTP commands. If the NVMe sStorage dDevice has an SMBus/I2C interfaceport, then the VPD is accessible using the access mechanism over I2C as defined in the IPMI Platform Management FRU Information Storage Definition.

Figure 2 illustrates an NVMe Storage Device that is a single-port PCIe SSD with the FRU Information Device implemented by the NVM Subsystem. Figure 3 illustrates an NVMe Storage Device that is a dual-port PCIe SSD with an SMBus/I2C port and a FRU Information Device implemented using a Serial EEPROM.

Figure 2: Single-Port PCIe SSD

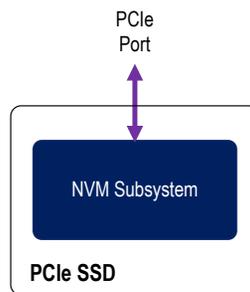
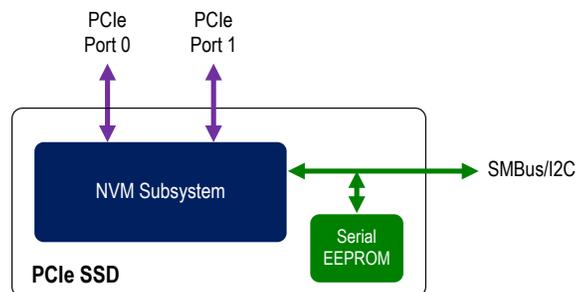


Figure 3: Dual-Port PCIe SSD with SMBus/I2C



The NVMe Management Interface is used to send Command Messages which consist of standard NVMe Admin Commands that target a Controller within the NVM Subsystem; commands that provide access to the PCI Express configuration, I/O, and memory spaces of a Controller in the NVM Subsystem; and Management Interface specific commands for inventorying, configuring and monitoring of the NVM Subsystem. Each Management Endpoints advertises its unique capabilities. All Management Endpoints

may support the same commands even though PCIe ports are full duplex with much higher data rates than SMBus (i.e., both SMBus/I2C and PCIe VDM are capable of providing the same functionality).

The PCIe ports and SMBus/I2C port of an NVM Subsystem may optionally each contain a single NVMe Management Endpoint (hereafter referred to as simply Management Endpoint). A Management Endpoint is an MCTP endpoint that is the terminus and origin of MCTP packets/messages and is responsible for implementing the MCTP Base Protocol, processing MCTP Control Messages, and internal routing of Command Messages.

Each NVMe Controller in the NVM Subsystem shall provide an NVMe Controller Management Interface (hereafter referred to as simply Controller Management Interface). The Controller Management Interface executes Controller operations on behalf of any Management Endpoint in the NVM Subsystem. Management Endpoints may route commands to any NVMe Controller in the NVM Subsystem. A Controller Management Interface logically executes one operation at a time. A Controller Management Interface is not precluded from executing two or more operations in parallel; however, there shall always be an equivalent pattern of sequential operations with the same results.

Figure 4 illustrates an example NVM Subsystem corresponding to the PCIe SSD shown in Figure 2. The NVM Subsystem contains a single Controller and there is a Management Endpoint associated with the PCIe port.

Figure 4: NVM Subsystem Associated with Single Ported PCIe SSD

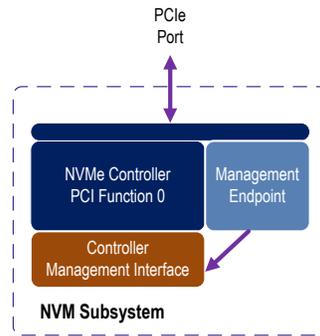
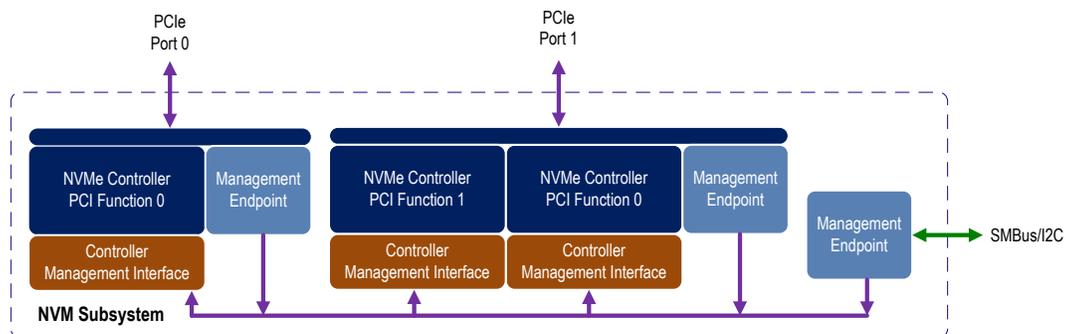


Figure 5 illustrates an example NVM Subsystem corresponding to the PCIe SSD shown in Figure 3. The NVM Subsystem contains one Controller associated with PCIe Port 0 and two Controllers associated with PCIe Port 1. There is a Management Endpoint associated with the each PCIe port and the SMBus/I2C port. Since the NVM Subsystem contains a Management Endpoint, all Controllers have an associated Controller Management Interface.

Figure 5: NVM Subsystem Associated with Dual Ported PCIe SSD with SMBus/I2C



Management Interface Request Messages and Response Messages are transported as MCTP messages with the Message Type set to NVM Express Management Messages over MCTP (refer to the MCTP IDs

and Codes specification). All Command Messages originate with the Management Controller and result in a Response Message from the Management Endpoint

Modify a portion of section 1.7 as follows:

NVM Express specification, revision 1.23. Available from <http://www.nvmexpress.org>

Modify a portion of figure 17 as follows:

21h	PCIe Inaccessible: The PCIe functionality is not available at this time.	Refer to 4.2.1
22h	Sanitize In Progress: The requested command is prohibited while a sanitize operation is in progress. Refer to section 9.1.	Refer to 4.2.1
223h – DFh	Reserved	

Update Section 6 Figure with the following:

The NVM Express Admin Command Set allows NVMe Admin commands to be issued to any Controller in the NVM Subsystem using NVMe-MI. Supported commands are listed in **Figure 76**, and are defined in the NVMe specification. If an NVMe Admin Command is issued in a Request Message **that is a prohibited command other than one listed in Figure 76**, the Management Endpoint shall return a response with status Invalid Parameter pointing to the NVMe opcode. Future revisions of this specification may add additional commands to **Figure 76**.

Figure 1: List of NVMe Admin Commands Supported

Command	O/M/P ¹
Abort	P
Asynchronous Event Request	P
Create I/O Completion Queue	P
Create I/O Submission Queue	P
Delete I/O Completion Queue	P
Delete I/O Submission Queue	P
Device Self-test	O
Directive Receive	P
Directive Send	P
Doorbell Buffer Config	P
Firmware Activate/Commit	O
Firmware Image Download	O
Format NVM	O
Get Features	M
Get Log Page	M
Identify	M
Keep Alive	P
Namespace Management	O
Namespace Attachment	O
NVMe-MI Receive	P
NVMe-MI Send	P
Sanitize	O
Security Send	O
Security Receive	O
Set Features	O
Vendor Specific	O
Virtualization Management	O
NOTES:	
1. O/M/P definition: O = Optional, M = Mandatory, P = Prohibited. Mandatory commands shall be supported if the NVMe Controller specified by the Controller ID field supports the command.	

Add a new section in after section 6.2 as follows:

6.2 Sanitize Operation

Figure X specifies the Command Messages allowed during a sanitize operation. Refer to the NVM Express specification for the definition of a sanitize operation.

Figure X: Command Messages Allowed During Sanitize Operation

Command Set	Command Message	Allowed During Sanitize Operation ¹
Management Interface Command Set	Configuration Get	Yes
	Configuration Set	
	Controller Health Status Poll	
	NVM Subsystem Health Status Poll	
	Read NVMe-MI Data Structure	
	Reset	
	SES Receive	
	SES Send	
	VPD Read	
	VPD Write	
NVMe Admin Command Set ²	Device Self-test	Same restrictions as defined by the NVM Express specification
	Firmware Activate/Commit	
	Firmware Image Download	
	Format NVM	
	Get Features	
	Get Log Page	
	Identify	
	Namespace Attachment	
	Namespace Management	
	Sanitize	
	Security ReceiveSend	
	Security Send	
	Set Features	
	Vendor Specific	
Virtualization Management		
PCIe Command Set	PCIe Configuration Read	Yes
	PCIe Configuration Write	
	PCIe I/O Read	
	PCIe I/O Write	
	PCIe Memory Read	
	PCIe Memory Write	
NOTES:		
1. Refer to the NVM Express specification for the definition of a sanitize operation.		
2. NVMe Admin Commands that are prohibited via the out-of-band mechanism (refer to Figure 76) are not listed since they are always prohibited including during a sanitize operation.		