



LEGAL NOTICE:

© Copyright 2007 - 2016 NVM Express, Inc. ALL RIGHTS RESERVED.

This erratum to the NVM Express Management Interface Revision 1.0 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 1.0 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 - 2016 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 Workgroup
c/o Virtual, Inc.
401 Edgewater Place, Suite 600
Wakefield, MA 01880
info@nvmexpress.org

NVM Express™ Technical Errata

Errata ID	003
Revision Date	9/12/2016
Affected Spec Ver.	NVM Express™ MI 1.0
Corrected Spec Ver.	

Errata Author(s)

Name	Company
Austin Bolen	Dell
Myron Loewen	Intel
Jinliang Mao	SK Hynix
John Wiedemeier	Teledyne LeCroy
Peter Onufryk	Microsemi

Errata Overview

The erratum includes clarifications for NVM Express Management Interface, Revision 1.0, including:

- Various editorial updates and clarifications.
- Many clarifications on the status bits and how they work.
- Correction that Total NVM Capacity is for the NVM Subsystem and not the Management Endpoint.
- VPD: Product Info Area is optional per IPMI FRU Data Storage Specification.
- VPD: Offsets of Product Info Area and Multi-Record Area are vendor specific.
- Clarified Response Data field size in Controller Health Status Poll.
- Response Entries is 1's based. Can't be 0's based since 0 is a valid value.
- Fixed one of the Message Integrity Check examples in Appendix B.
- Reset NVM Subsystem is only required if it supported in-band
- Clarified that Invalid Parameter should be returned Request Messages with non-zero reserved fields
- Corrected "Invalid Field" and "Invalid Field in Command" to "Invalid Parameter"

Revision History

Revision Date	Change Description
7/1/2016	Initial draft
9/12/2016	Final revision

Description of Specification Changes

Modify a Figure 44 as shown below:

Figure 44: MCTP Transmission Unit Size – NVMe Management Response

Bit	Description
23:1604	Reserved
15:00	MCTP Transmission Unit Size: This field contains the MCTP Transmission Unit Size in bytes to be used by the port. The default value for this field following a reset or power cycle is 40h (64).

Modify a portion of Figure 76: NVMe Admin Command Request Description as shown below:

07:06	Controller ID (CTLID): This field specifies the C on t roller ID of the NVMe -Controller that this command targets.
-------	---

Modify a portion of table in Section 9.2.3 (NVMe MultiRecord Area) as shown below:

36:24	Impl Spec	Total NVM Capacity: This field indicates the total NVM cap <u>a</u> city of the Management Endpoint <u>NVM Subsystem</u> in bytes. If the NVM Subsystem supports Namespace Management, then this field should correspond to the value reported in the TNVMCAP field in the NVMe Identify Controller Data structure. A value of 0h may be used to indicate this feature is not supported.
-------	--------------	--

Modify Figure 108 (VPD Elements) as shown below:

Figure 108: VPD Elements

Byte	Name
07:00	Common Header
Vendor Specific 119:08	Product Info Area <u>(optional)</u>
Vendor Specific:120	MultiRecord Info Area
Vendor Specific	Internal Use Area (optional)

Vendor Specific	Chassis Info Area (optional)
Vendor Specific	Board Info Area (optional)

Modify portion of Figure 47 (SMBus/I2C Frequency – NVMe Management Dword 0) as shown below:

110:08	SMBus/I2C Frequency: This field specifies the new frequency for the specified SMBus/I2C port.	
	Value	Description
	0h	Reserved
	1h	100 kHz
	2h	400 kHz
	3h	1 MHz
	4h - Fh	Reserved

Modify the Byte column in Figures 27, 28, 29, 30, and 31 for the Control Primitive Specific Response (CPSR) field as shown below:

Byte
078:067

Modify Figure 32 (Replay Control Primitive Success Response Message Fields) as shown below:

Byte	Description
------	-------------

078:067	Control Primitive Specific Response (CPSR): This field is used to return Control Primitive specific status.	
	Bit	Description
	15:01	Reserved
	00	Response Replay (RR): This bit indicates if a previous Response Message is retransmitted. This field is set to '1' if the requested Response Message is retransmitted by the Management Endpoint. This field is cleared to '0' if the requested Response Message is not retransmitted.

Make the following modifications to Figure 30:

Note to technical writer: The table in Figure 30 is formatted incorrectly. The figure caption is on page 36 followed by a large gap. The first row of the table is the only thing on page 37. The rest of the table starts on page 38. Need to group the figure caption and first row of the table with the rest of the table.

Modify a portion of Figure 49 (Health Status Change – NVMe Management Dword 1) as shown below:

Bit	Description
31:12	Reserved
11	Critical Warning (CWARN): When this bit is set to '1', the corresponding bit 12 in the Composite Controller Status field of the NVM Subsystem Health Data Structure is cleared to '0'.
10	Available Spare (SPARE): When this bit is set to '1', the corresponding bit 11 in the Composite Controller Status field of the NVM Subsystem Health Data Structure is cleared to '0'.
9	Percentage Used (PDLU): When this bit is set to '1', the corresponding bit 10 in the Composite Controller Status field of the NVM Subsystem Health Data Structure is cleared to '0'.
8	Composite Temperature Change (CTEMP): When this bit is set to '1', the corresponding bit 9 in the Composite Controller Status field of the NVM Subsystem Health Data Structure is cleared to '0'.
7	Controller Status Change: When this bit is set to '1', the corresponding bit 8 in the Composite Controller Status field of the NVM Subsystem Health Data Structure is cleared to '0'.
6	Firmware Activated (FA): When this bit is set to '1', the corresponding bit 7 in the Composite Controller Status field of the NVM Subsystem Health Data Structure is cleared to '0'.
5	Namespace Attribute Changed (NAC): When this bit is set to '1', the corresponding bit 6 in the Composite Controller Status field of the NVM Subsystem Health Data Structure is cleared to '0'.
4	Controller Enable Change Occurred (CECO): When this bit is set to '1', the corresponding bit 5 in the Composite Controller Status field of the NVM Subsystem Health Data Structure is cleared to '0'.
3	NVM Subsystem Reset Occurred (NSSRO): When this bit is set to '1', the corresponding bit 4 in the Composite Controller Status field of the NVM Subsystem Health Data Structure is cleared to '0'.

2	Shutdown Status (SHST): When this bit is set to '1', the corresponding bit 2 in the Composite Controller Status field of the NVM Subsystem Health Data Structure is cleared to '0'.
1	Controller Fatal Status (CFS): When this bit is set to '1', the corresponding bit 1 in the Composite Controller Status field of the NVM Subsystem Health Data Structure is cleared to '0'.
0	Ready (RDY): When this bit is set to '1', the corresponding bit 0 in the Composite Controller Status field of the NVM Subsystem Health Data Structure is cleared to '0'.

Modify a portion of Figure 52 (Controller Health Status Poll – NVMe Management Dword 0) as shown below:

Bit	Description
31	Report All (ALL): When this bit is set to '1', health status <u>a Controller Health Data Structure</u> is returned for Controllers regardless of the status of the <u>Controller Health Status Changed</u> Flags bit vector (i.e., it is as though all the bits are set in the Controller Health Status Changed <u>Flags bit vector</u>).
30:27	Reserved
26	Include SR-IOV Virtual Functions (INCVF): When this bit is set to 1, a Controller Health Data Structure is returned <u>Controller Health Status is reported</u> for NVMe-Controllers associated with SR-IOV Virtual Functions (VFs).
25	Include SR-IOV Physical Functions (INCPF): When this bit is set to 1, a Controller Health Data Structure is returned <u>Controller Health Status is reported</u> for NVMe-Controllers associated with SR-IOV Physical Functions (PFs).
24	Include PCI Functions (INCF): When this bit is set to 1, a Controller Health Data Structure is returned <u>Controller Health Status is reported</u> for NVMe-Controllers associated with a non SR-IOV PCI Function.
23:16	Maximum Response Entries (MAXRENT): This field specifies the maximum number of Controller Health Data Structure entries that may be returned in the completion. This is <u>a</u> 0's based field. The maximum number of entries is 255. Specifying 256 entries is interpreted as an Invalid <u>Field Parameter</u> .
15:00	Starting Controller ID (SCTLID): This field specifies the starting Controller ID from which to return of the first Controller whose <u>Controller Health Data Structure may be returned</u> health status information.

Modify a portion of Figure 53 (Controller Health Status Poll – NVMe Management Dword 1) as shown below:

Bit	Description
31	Clear Changed Flags (CCF): When this bit is set to 1, the <u>Controller Health Status Changed</u> Flags state of reported changed flag bits in the changed flag bit vector are cleared in Controllers whose health status <u>Controller Health Data Structure</u> is contained in the Response Data.
30:5	Reserved

04	Critical Warning (CWARN): When this bit is set to 1, <u>a Controller Health Data Structure is returned for Controllers with the Critical Warning bit set in their Controller Health Status Changed Flagscritical-warning-changes are reported.</u>
03	Available Spare (SPARE): When this bit is set to 1, <u>a Controller Health Data Structure is returned for Controllers with the Available Spare bit set in their Controller Health Status Changed Flagsavailable-spare-changes are reported.</u>
02	Percentage Used (PDLU): When this bit is set to 1, <u>a Controller Health Data Structure is returned for Controllers with the Percent Used bit set in their Controller Health Status Changed Flagspercentage-used-changes are reported.</u>
01	Composite Temperature Changes (CTEMP): When this bit is set to 1, <u>a Controller Health Data Structure is returned for Controllers with the Composite Temperature bit set in their Controller Health Status Changed Flagsecomposite-temperature-changes are reported.</u>
00	Controller Status Changes (CSTS): When this bit is set to 1, <u>a Controller Health Data Structure is returned for Controllers with the Ready, Controller Fatal Status, Shutdown Status, NVM Subsystem Reset Occurred, Controller Enable Change Occurred, Namespace Attribute Changed, or Firmware Activated bit set in their Controller Health Status Changed FlagsController status-changes are reported.</u>

Modify a portion of Section 5.3 (Controller Health Status Poll) as shown below:

Associated with each Controller in the NVM Subsystem is a set of Controller Health Status Changed Flagsflag bit-vector shown in Figure XX1. The Controller Health Status Changed Flags are set when the corresponding field in the Controller Health Data Structure changes state as described in Figure XX1. Figure XX2 shows a graphical representation of which field(s)/bit(s) in the Controller Health Data Structure are associated with each bit in the Controller Health Status Changed Flags. When a bit in the Controller Health Status Changed Flags for any Controller transitions from '0' to '1', then the corresponding bit in the Composite Controller Status is also set to '1', with a bit corresponding to each field in the Controller Health Data Structure. The initial value following a reset or power cycle of all changed flag bits is cleared to '0'. A Health Status Changed flag bit in the bit-vector for a Controller is set when the value of the corresponding Controller Health Data Structure field for that Controller changes state. The state of the entire changed flag bit vector is cleared in a Controller on a reset, power cycle, or a Configuration Set command that selects Health Status Change. The Controller Health Status Changed Flags state of reported bits in the changed flag bit vector isare cleared in Controllers whose health statusController Health Data Structure is returned in the Success Response Message to a Controller Health Status Poll ~~e~~Command Message with the Clear Changed Flags bit set to '1'.

A Controller Health Status Poll response may return the ~~health status~~Controller Health Data Structure for up to 255 Controllers in the Response Data field. An NVM Subsystem may contain up to 64K Controllers, so a method is needed to limit the size of the Response Message. The Starting Controller ID~~entifier~~ field in the Command Message specifies the ~~starting~~ Controller ID of the first Controller whose Controller Health Data Structure may be returned in the Response Data fieldthat is checked for a change in health status, while ~~t~~The Maximum Response Entries field specifies the maximum number of Controllers whose Controller Health Data Structurehealth status may be returned in the Response Data field.

-The Response Data field contains the Controller Health Status Data Structure for up to the first M Controllers starting with Controller N ~~whose health status has changed~~, where M is equal to the Maximum Response Entries field and N is equal to the Starting Controller ID~~entifier~~ field. The Response Data field shall contain the Controller Health Status Data Structure for all Controllers that do not match the filtering criteria in Controller Health Status Poll - NVMe Management Dword 0 (refer to Section 5.3.1) and that have one or more Controller Health Status Changed Flags that are a) set and b) do not match the filtering criteria in Controller Health Status Poll - NVMe Management Dword 1 (refer to Section 5.3.2). The Response Data field shall not contain the Controller Health Status Data Structure for any Controllers that meet the filtering criteria in Sections 5.3.1 or 5.3.2.

5.3.1 Filtering by Controller Type

The Controller Health Data Structures that are returned by Controller Health Status Poll ~~Health-status~~ may be filtered (i.e., excluded from being included in the Response Data field regardless of the state of the Controller Health Status Changed Flags) by ~~NVMe~~ Controller type (i.e., non SR-IOV PCI Function, SR-IOV PF, and SR-IOV VF) ~~and by fields in the Controller Health Data Structure~~. ~~NVMe~~ Controller type filtering is selected controlled by the Include PCI Functions, Include SR-IOV PFs, and Include SR-IOV VFs fields in NVMe Management Dword 0. When one of these bits is set, Controller Health Data Structures for Controllers corresponding to that type of PCI Function are ~~excluded~~included in the Response Data field; else, the Controller Health Data Structure for that Controller is excluded from the Response Data field.

5.3.2 Filtering by Controller Health Status Changed Flags

The Controller Health Data Structures that are returned by Controller Health Status Poll may also be filtered by the Controller Health Status Changed Flags. Filtering of changes by Controller Health Status Changed Flags ~~in individual Controller Health Data Structure fields~~ is controlled by ~~fields~~ some of the bits in NVMe Management Dword 1. When one or more of these bits is set and any of the corresponding bit(s) in the Controller Health Status Changed Flags for the Controller are also set (refer to Figure 53 for Controller Health Status Changed Flags associated with each bit in NVMe Management Dword 1), then the entire Controller Health Data Structure (including any filtered fields) for that Controller is returned in the Response Data field; else, the Controller Health Data Structure for that Controller is excluded from the Response Data field. The contents returned in the Controller Health Data Structure for filtered fields are undefined. ~~When one of these bits is cleared to '0', then that field is removed from determination of health status changes for that Controller. A Controller's health status is considered to have changed when one or more unfiltered changed flag bits in the bit vector for that Controller are set.~~

Figure XX1 Controller Health Status Changed Flags (CHSCF)

Bit	Reset	Description
15:13	0	Reserved
12	0	Critical Warning (CWARN): This bit is set to '1' when any of the Critical Warning bits in the Controller Health Data Structure transition from '0' to '1'.
11	0	Available Spare (SPARE): This bit is set to '1' when the Available Spare field in the Controller Health Data Structure changes state.
10	0	Percentage Used (PDLU): This bit is set to '1' when the Percentage Used field in the Controller Health Data Structure changes state.
9	0	Composite Temperature Change (CTEMP): This bit is set to '1' when the Composite Temperature field in the Controller Health Data Structure changes state.
8	Hwlnit	Controller Status Change (CSTS): This bit is set to '1' when the Shutdown Status field in the Controller Health Data Structure changes state or when the Ready, Controller Fatal Status, NVM Subsystem Reset Occurred, Controller Enable Change Occurred, Namespace Attribute Changed, or Firmware Activated bit in the in the Controller Health Data Structure transitions from '0' to '1'.
7	Hwlnit	Firmware Activated (FA): This bit is set to '1' when the Firmware Activated bit in the Controller Health Data Structure transitions from '0' to '1'.
6	0	Namespace Attribute Changed (NAC): This bit is set to '1' when the Namespace Attribute Changed bit in the Controller Health Data Structure transitions from '0' to '1'.
5	0	Controller Enable Change Occurred (CECO): This bit is set to '1' when the Controller Enable Change Occurred bit in the Controller Health Data Structure transitions from '0' to '1'.
4	Hwlnit	NVM Subsystem Reset Occurred (NSSRO): This bit is set to '1' when the NVM Subsystem Reset Occurred bit in the Controller Health Data Structure transitions from '0' to '1'.
3	0	Reserved
2	0	Shutdown Status (SHST): This bit is set to '1' when the Shutdown Status field in the Controller Health Data Structure changes state.
1	Hwlnit	Controller Fatal Status (CFS): This bit is set to '1' when the Controller Fatal Status bit in the Controller Health Data Structure transitions from '0' to '1'.
0	0	Ready (RDY): This bit is set to '1' when the Ready bit in the Controller Health Data Structure transitions from '0' to '1'.

Figure XX2 Controller Health Data Structure to Controller Health Status Changed Flags Mapping

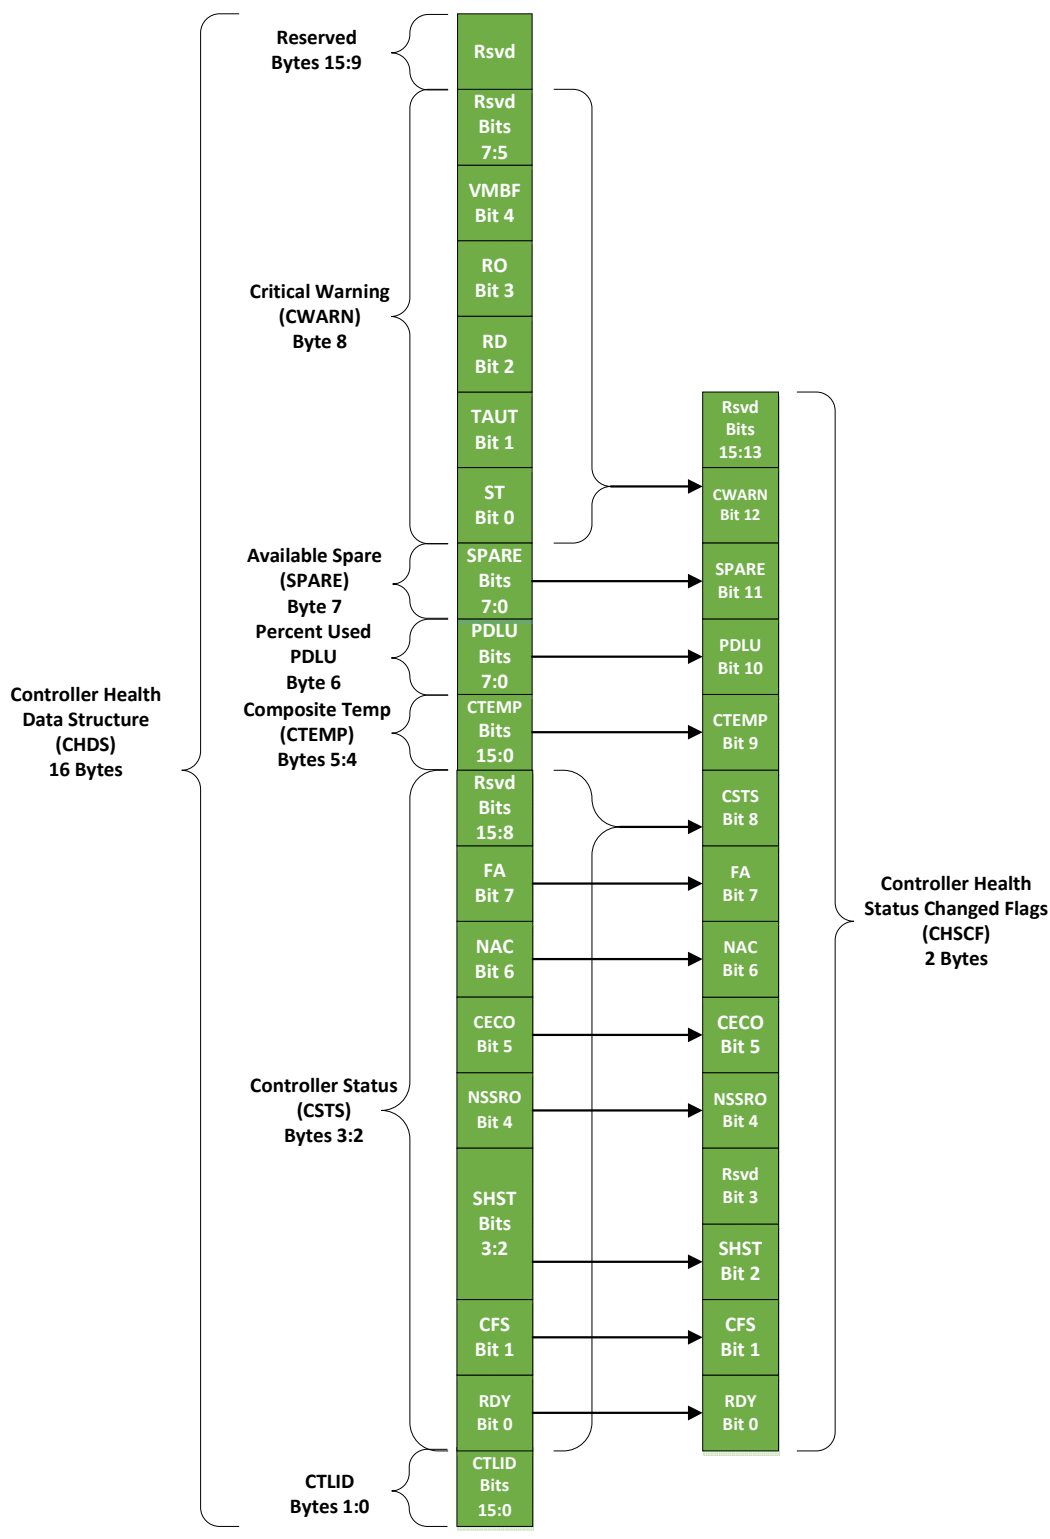


Figure 55: Controller Health Data Structure (CHDS)

Bytes	Description																											
1:0	Controller Identifier (CTLID): This field specifies the NVMe Controller Identifier with which the data contained in this data structure is associated.																											
3:2	Controller Status (CSTS): This field reports the NVMe Controller status. <table><tr><th>Bit</th><th>Reset</th><th>Description</th></tr><tr><td>15:8</td><td><u>0</u></td><td>Reserved</td></tr><tr><td>7</td><td><u>HwInit</u></td><td>Firmware Activated (FA): This bit is set to '1' when a new firmware image is activated. This bit is cleared to '0' after it is read using this command. <u>The reset value of this field is '1' if a reset caused a new firmware image to be activated.</u></td></tr><tr><td>6</td><td><u>0</u></td><td>Namespace Attribute Changed (NAC): This bit is set to '1' when a change occurs in the Identify Namespace data structure for one or more namespaces. This bit is cleared to '0' after it is read using this command.</td></tr><tr><td>5</td><td><u>0</u></td><td>Controller Enable Change Occurred (CECO): This bit is set to '1' when the Enable (CC.EN) bit (<u>refer to CC.EN in the NVM Express specification</u>) changes state. This bit is cleared to '0' after it is read using this command.</td></tr><tr><td>4</td><td><u>HwInit</u></td><td>NVM Subsystem Reset Occurred (NSSRO): This bit corresponds to the value of the NVM Subsystem Reset Occurred (<u>refer to CSTS.NSSRO in the NVM Express specification</u>) bit.</td></tr><tr><td>3:2</td><td><u>0</u></td><td>Shutdown Status (SHST): This field corresponds to the value of the Shutdown Status (<u>refer to CSTS.SHST in the NVM Express specification</u>) field.</td></tr><tr><td>1</td><td><u>HwInit</u></td><td>Controller Fatal Status (CFS): This bit corresponds to the value of the Controller Fatal Status (<u>refer to CSTS.CFS in the NVM Express specification</u>) bit.</td></tr><tr><td>0</td><td><u>0</u></td><td>Ready (RDY): This bit corresponds to the value of the Ready (<u>refer to CSTS.RDY in the NVM Express specification</u>) bit.</td></tr></table>	Bit	Reset	Description	15:8	<u>0</u>	Reserved	7	<u>HwInit</u>	Firmware Activated (FA): This bit is set to '1' when a new firmware image is activated. This bit is cleared to '0' after it is read using this command. <u>The reset value of this field is '1' if a reset caused a new firmware image to be activated.</u>	6	<u>0</u>	Namespace Attribute Changed (NAC): This bit is set to '1' when a change occurs in the Identify Namespace data structure for one or more namespaces. This bit is cleared to '0' after it is read using this command.	5	<u>0</u>	Controller Enable Change Occurred (CECO): This bit is set to '1' when the Enable (CC.EN) bit (<u>refer to CC.EN in the NVM Express specification</u>) changes state. This bit is cleared to '0' after it is read using this command.	4	<u>HwInit</u>	NVM Subsystem Reset Occurred (NSSRO): This bit corresponds to the value of the NVM Subsystem Reset Occurred (<u>refer to CSTS.NSSRO in the NVM Express specification</u>) bit.	3:2	<u>0</u>	Shutdown Status (SHST): This field corresponds to the value of the Shutdown Status (<u>refer to CSTS.SHST in the NVM Express specification</u>) field.	1	<u>HwInit</u>	Controller Fatal Status (CFS): This bit corresponds to the value of the Controller Fatal Status (<u>refer to CSTS.CFS in the NVM Express specification</u>) bit.	0	<u>0</u>	Ready (RDY): This bit corresponds to the value of the Ready (<u>refer to CSTS.RDY in the NVM Express specification</u>) bit.
	Bit	Reset	Description																									
	15:8	<u>0</u>	Reserved																									
	7	<u>HwInit</u>	Firmware Activated (FA): This bit is set to '1' when a new firmware image is activated. This bit is cleared to '0' after it is read using this command. <u>The reset value of this field is '1' if a reset caused a new firmware image to be activated.</u>																									
	6	<u>0</u>	Namespace Attribute Changed (NAC): This bit is set to '1' when a change occurs in the Identify Namespace data structure for one or more namespaces. This bit is cleared to '0' after it is read using this command.																									
	5	<u>0</u>	Controller Enable Change Occurred (CECO): This bit is set to '1' when the Enable (CC.EN) bit (<u>refer to CC.EN in the NVM Express specification</u>) changes state. This bit is cleared to '0' after it is read using this command.																									
	4	<u>HwInit</u>	NVM Subsystem Reset Occurred (NSSRO): This bit corresponds to the value of the NVM Subsystem Reset Occurred (<u>refer to CSTS.NSSRO in the NVM Express specification</u>) bit.																									
	3:2	<u>0</u>	Shutdown Status (SHST): This field corresponds to the value of the Shutdown Status (<u>refer to CSTS.SHST in the NVM Express specification</u>) field.																									
	1	<u>HwInit</u>	Controller Fatal Status (CFS): This bit corresponds to the value of the Controller Fatal Status (<u>refer to CSTS.CFS in the NVM Express specification</u>) bit.																									
0	<u>0</u>	Ready (RDY): This bit corresponds to the value of the Ready (<u>refer to CSTS.RDY in the NVM Express specification</u>) bit.																										
5:4	Composite Temperature (CTEMP): This field contains a value corresponding to a temperature in degrees Kelvin that represents the current composite temperature of the Controller and namespace(s) associated with that Controller. The value of this field corresponds to the value in the NVMe Controller's SMART / Health Information Log.																											
6	Percentage Used (PDLU): This field contains a vendor specific estimate of the percentage of NVM Subsystem life used based on the actual usage and the manufacturer's prediction of NVM life. The value of this field corresponds to the value in the NVMe Controller's SMART / Health Information Log.																											
7	Available Spare (SPARE): This field contains a normalized percentage (0 to 100%) of the remaining spare capacity available. The value of this field corresponds to the value in the NVMe Controller's SMART / Health Information Log.																											

8	Critical Warning (CWARN): This field indicates critical warnings for the state of the Controller. The value of this field corresponds to the value in the NVMe Controller's SMART / Health Information Log.	
	Bit	Description
	7:5	Reserved
	4	Volatile Memory Backup Failed (VMBF): This bit is set to '1' when the volatile memory backup device has failed.
	3	Read Only (RO): This bit is set to '1' when the media has been placed in read only mode.
	2	Reliability Degraded (RD): This bit is set to '1' when NVM Subsystem reliability has been degraded due to significant media related errors or an internal error.
	1	Temperature Above or Under Threshold (TAUT): This bit is set to '1' when a temperature is above an over temperature threshold or below an under temperature threshold.
	0	Spare Threshold (ST): this bit is set to '1' when the available spare has fallen below the available spare threshold.
15:9	Reserved	

Modify a portion of Figure 57 (NVM Subsystem Health Data Structure) as shown below:

Figure 57: NVM Subsystem Health Data Structure **(NSHDS)**

Byte	Description														
0	<p>NVM Subsystem Status <u>(NSS)</u>: This field indicates the status of the NVM Subsystem.</p> <table> <tr> <th>Bit</th><th>Description</th></tr> <tr> <td>7:6</td><td>Reserved</td></tr> <tr> <td>5</td><td>Drive Functional <u>(DF)</u>: This bit is set to '1' to indicate an NVM Subsystem is functional. If cleared to '0', then there is an unrecoverable failure <u>detected</u> in the NVM Subsystem and the rest of the transmission may contain invalid information.</td></tr> <tr> <td>4</td><td>Reset Not Required <u>(RNR)</u>: This bit is set to '1' to indicate the NVM Subsystem does not need a reset to resume normal operation. If cleared to '0' then the NVM Subsystem has experienced an error that prevents continued normal operation. A Controller Level Reset is required to resume normal operation.</td></tr> <tr> <td>3</td><td>Port 0 PCIe Link Active <u>(P0LA)</u>: This bit is set to '1' to indicate the first port's PCIe link is up (i.e., the Data Link Control and Management State Machine is in the DL_Active state). If cleared to '0', then the PCIe link is down.</td></tr> <tr> <td>2</td><td>Port 1 PCIe Link Active <u>(P1LA)</u>: This bit is set to '1' to indicate the second port's PCIe link is up. If cleared to '0', then the second port's PCIe link is down or not present.</td></tr> <tr> <td>1:0</td><td>Reserved</td></tr> </table>	Bit	Description	7:6	Reserved	5	Drive Functional <u>(DF)</u>: This bit is set to '1' to indicate an NVM Subsystem is functional. If cleared to '0', then there is an unrecoverable failure <u>detected</u> in the NVM Subsystem and the rest of the transmission may contain invalid information.	4	Reset Not Required <u>(RNR)</u>: This bit is set to '1' to indicate the NVM Subsystem does not need a reset to resume normal operation. If cleared to '0' then the NVM Subsystem has experienced an error that prevents continued normal operation. A Controller Level Reset is required to resume normal operation.	3	Port 0 PCIe Link Active <u>(P0LA)</u>: This bit is set to '1' to indicate the first port's PCIe link is up (i.e., the Data Link Control and Management State Machine is in the DL_Active state). If cleared to '0', then the PCIe link is down.	2	Port 1 PCIe Link Active <u>(P1LA)</u>: This bit is set to '1' to indicate the second port's PCIe link is up. If cleared to '0', then the second port's PCIe link is down or not present.	1:0	Reserved
Bit	Description														
7:6	Reserved														
5	Drive Functional <u>(DF)</u>: This bit is set to '1' to indicate an NVM Subsystem is functional. If cleared to '0', then there is an unrecoverable failure <u>detected</u> in the NVM Subsystem and the rest of the transmission may contain invalid information.														
4	Reset Not Required <u>(RNR)</u>: This bit is set to '1' to indicate the NVM Subsystem does not need a reset to resume normal operation. If cleared to '0' then the NVM Subsystem has experienced an error that prevents continued normal operation. A Controller Level Reset is required to resume normal operation.														
3	Port 0 PCIe Link Active <u>(P0LA)</u>: This bit is set to '1' to indicate the first port's PCIe link is up (i.e., the Data Link Control and Management State Machine is in the DL_Active state). If cleared to '0', then the PCIe link is down.														
2	Port 1 PCIe Link Active <u>(P1LA)</u>: This bit is set to '1' to indicate the second port's PCIe link is up. If cleared to '0', then the second port's PCIe link is down or not present.														
1:0	Reserved														
1	<p>Smart Warnings <u>(SW)</u>: This field contains the Critical Warning field (byte 0) of the NVMe SMART / Health Information log. Each bit in this field is inverted from the NVMe definition (i.e., the management interface shall indicate a '0' value while the corresponding bit is '1' in the log page). Refer to the NVMe specification for bit definitions.</p> <p>If there are multiple Controllers in the NVM Subsystem, the management endpoint shall combine the Critical Warning field from every Controller in the NVM Subsystem such that a bit in this field is:</p> <ul style="list-style-type: none"> Cleared to '0' if any Controller in the subsystem indicates a critical warning for that corresponding bit. Set to '1' if all Controllers in the NVM Subsystem do not indicate a critical warning for the corresponding bit. 														
2	<p>Composite Temperature <u>(CTEMP)</u>: This field indicates the current temperature in degrees Celsius. If a temperature value is reported, it should be the same temperature as the Composite Temperature from the SMART log of hottest Controller in the NVM Subsystem. The reported temperature range is vendor specific, and shall not exceed the range -60 to +127°C. The 8 bit format of the data is shown below.</p> <p>This field should not report a temperature that is older than 1 second. If recent data is not available, the Management Endpoint should indicate a value of 80h for this field.</p> <table> <tr> <th>Value</th><th>Description</th></tr> <tr> <td>00h-7Eh</td><td>Temperature is measured in degrees Celsius (0 to 126C)</td></tr> <tr> <td>7Fh</td><td>127C or higher</td></tr> <tr> <td>80h</td><td>No temperature data or temperature data is more the 5 seconds old.</td></tr> <tr> <td>81h</td><td>Temperature sensor failure</td></tr> </table>	Value	Description	00h-7Eh	Temperature is measured in degrees Celsius (0 to 126C)	7Fh	127C or higher	80h	No temperature data or temperature data is more the 5 seconds old.	81h	Temperature sensor failure				
Value	Description														
00h-7Eh	Temperature is measured in degrees Celsius (0 to 126C)														
7Fh	127C or higher														
80h	No temperature data or temperature data is more the 5 seconds old.														
81h	Temperature sensor failure														

		82h-C3h	Reserved	
		C4	Temperature is -60C or lower	
		C5-FFh	Temperature measured in degrees Celsius is represented in two's complement (-1 to -59C)	
3	Percentage Drive Life Used (PDLU): Contains a vendor specific estimate of the percentage of NVM Subsystem NVM life used based on the actual usage and the manufacturer's prediction of NVM life. If an NVM Subsystem has multiple Controllers the highest value is returned. A value of 100 indicates that the estimated endurance of the NVM in the NVM Subsystem has been consumed, but may not indicate an NVM Subsystem failure. The value is allowed to exceed 100. Percentages greater than 254 shall be represented as 255. This value should be updated once per power-on hour and equal the Percentage Used value in the NVMe SMART Health Log Page.			
5:4	Composite Controller Status (CCS): This field reports the composite status of all Controllers in the NVM Subsystem.			
	All bits in this field are cleared to '0' for a Controller during a Controller Level Reset. The bits in this field are also cleared after the NVM Subsystem Health Data Structure (refer to Figure 57) is returned in a Success Response Message associated with a NVM Subsystem Health Status Poll command where the Clear Status bit set. A Configuration Set command that selects Health Status Change may be used to clear selected bits to '0'.			
	Bit	Reset	Description	
	15:13	0	Reserved	
	12	0	Critical Warning (CWARN): This bit is set to '1' when the Critical Warning field- bit in the Controller Health Data-Structure Status Changed Flags is set transitions from '0' to '1' in one or more Controllers in the NVM Subsystem.	
	11	0	Available Spare (SPARE): This bit is set to '1' when the Available Spare bitfield in the Controller Health Status Changed Flags Data-Structure transitions from '0' to '1' has changed state in one or more Controllers in the NVM Subsystem.	
	10	0	Percentage Used (PDLU): This bit is set to '1' when the Percentage Used bitfield in the Controller Health Status Changed Flags Data-Structure transitions from '0' is set to '1' in one or more Controllers in the NVM Subsystem.	
	9	0	Composite Temperature Change (CTEMP): This bit is set to '1' when the Composite Temperature bitfield in the Controller Health Status Changed Flags Data-Structure transitions from '0' is set to '1' in one or more Controllers in the NVM Subsystem.	
	8	HwInit	Controller Status Change (CSTS): This bit is set to '1' when the Controller Status field in the Controller Health Status Changed Flags Data-Structure transitions from '0' is set to '1' in one or more Controllers in the NVM Subsystem.	
	7	HwInit	Firmware Activated (FA): This bit is set to '1' when the Firmware Activated bit in the Controller Health Status Changed Flags transitions from '0' to '1' in one or more Controllers in the NVM Subsystem. This bit is set to '1' when a new firmware image is activated in the NVM Subsystem.	
	6	0	Namespace Attribute Changed (NAC): This bit is set to '1' when the	

			<u>Namespace Attribute Changed bit in the Controller Health Status Changed Flags transitions from '0' to '1' in one or more Controllers in the NVM Subsystem. This bit is set to '1' when a change occurred in the Identify Namespace data structure associated with one or more namespaces in the NVM Subsystem.</u>
5	<u>0</u>		Controller Enable Change Occurred (CECO): This bit is set to '1' when the Controller Enable Change Occurred bit in the Controller Health Status Changed Flags transitions from '0' to '1' in one or more Controllers in the NVM Subsystem. This bit is set to '1' when the Enable (CC.EN) bit changes state (i.e., '0' to '1' or '1' to '0') in one or more Controllers in the NVM Subsystem.
4	<u>HwInit</u>		NVM Subsystem Reset Occurred (NSSRO): This bit is set to '1' when the NVM Subsystem Reset Occurred bit in the Controller Health Status Changed Flags transitions from '0' to '1' in one or more Controllers in the NVM Subsystem. This bit is set to '1' when the value of the NVM Subsystem Reset Occurred (CSTS.NSSRO) bit transitions from a '0' to a '1' in one or more Controllers in the NVM Subsystem.
<u>3</u>	<u>0</u>		Reserved
<u>3:2</u>	<u>0</u>		Shutdown Status (SHST): This bit is set to '1' when the Shutdown Status bit in the Controller Health Status Changed Flags transitions from '0' to '1' in one or more Controllers in the NVM Subsystem. This bit is set to '1' when the value of the Shutdown Status (CSTS.SHST) field bit transitions from a '0' to a '1' in one or more Controllers in the NVM Subsystem.
1	<u>HwInit</u>		Controller Fatal Status (CFS): This bit is set to '1' when the Controller Fatal Status bit in the Controller Health Status Changed Flags transitions from '0' to '1' in one or more Controllers in the NVM Subsystem. This bit is set to '1' when the value of the Controller Fatal Status (CSTS.CFS) bit transitions from a '0' to a '1' in one or more Controllers in the NVM Subsystem.
0	<u>0</u>		Ready (RDY): This bit is set to '1' when the Ready bit in the Controller Health Status Changed Flags transitions from '0' to '1' in one or more Controllers in the NVM Subsystem. This bit is set to '1' when the value of the Ready (CSTS.RDY) bit transitions from a '0' to a '1' in one or more Controllers in the NVM Subsystem.

Modify a portion of section 5.3 (Controller Health Status Poll) as shown below:

The Response Data field size may vary based on the number of Controllers whose health status-Controller Health Data Structure has changed and which-based on the number of Controllers whose Controller Health Data Structure isfields and Controller types are filtered _out by Controller type (refer to Section 5.3.1) or Controller Health Status Changed Flags (refer to Section 5.3.2).

Modify Figure 54 (Controller Health Status Poll – NVMe Management Response) as shown below:

Bit	Description
23:16	Response Entries (RENT): This field specifies the number of Controller Health Data Structure Entries present in the Response Data for this Response Message. This is a 0's based value.
15:00	Reserved

Modify a portion of Section 9.3.1 (NVM Subsystem Reset) as follows:

When an NVM Subsystem Reset is initiated, the entire NVM Subsystem is reset. This includes all NVM Subsystem ports (PCIe and SMBus/I2C), Management Endpoints, and Controller Management Interfaces. All state is returned to its default condition.

Modify Figure 114 (MIC Example 4 – 32 Decrementing Bytes from 0x1F to 0x00) as shown below:

	3	2	1	0
Dword 0	1Ch	1Dh	1Eh	1Fh
...
Dword 7	003h	012h	024h	093h
Dword 8 (MIC)	11h	3Fh	DBh	5Ch

Modify Figure 35: Opcodes for Management Interface Commands as shown below:

Opcode	O/M ¹	Command
00h	M	Read NVMe-MI Data Structure
01h	M	NVM Subsystem Health Status Poll
02h	M	Controller Health Status Poll
03h	M	Configuration Set
04h	M	Configuration Get
05h	M	VPD Read
06h	M	VPD Write
07h	OM	Reset
08h – BFh	-	Reserved
C0h – FFh	O	Vendor specific
NOTES:		
1. O/M definition: O = Optional, M = Mandatory.		

Modify a portion of Section 5.6 (Reset) as shown below:

Figure 67: Reset - NVMe Management Dword 0

Bit	Description									
31:24	Reset Type: This field specifies the type of reset to be performed.									
	<table><tr><th>Value</th><th><u>O/M¹</u></th><th>Description</th></tr><tr><td>00h</td><td><u>O/M²</u></td><td>Reset NVM Subsystem</td></tr><tr><td>01h – FFh</td><td>-</td><td>Reserved</td></tr></table>	Value	<u>O/M¹</u>	Description	00h	<u>O/M²</u>	Reset NVM Subsystem	01h – FFh	-	Reserved
	Value	<u>O/M¹</u>	Description							
	00h	<u>O/M²</u>	Reset NVM Subsystem							
01h – FFh	-	Reserved								
23:00	Reserved									
<u>NOTES:</u> 1. <u>O/M definition: O = Optional, M = Mandatory.</u> 2. <u>The Reset Type for Reset NVM Subsystem is required if the NVM Subsystem Reset feature is supported in-band as defined in the NVM Express specification; else, it is optional.</u>										

Modify a portion of Figure 110: Subsystem Management Data Structure as shown below:

Command Code	Offset (byte)	Description
0	00	Length of Status: Indicates number of additional bytes to read before encountering PEC. This value should always be 6 (06h) in implementations of this version of the spec.

Command Code	Offset (byte)	Description
	01	<p>Status Flags (SFLGS): This field indicates the status of the NVM Subsystem.</p> <p>SMBus Arbitration – Bit 7 is set ‘1’ after an SMBus block read is completed all the way to the stop bit without bus contention and cleared to ‘0’ if an SMBus Send Byte FFh is received on this SMBus slave address.</p> <p>Drive Not Ready – Bit 6 is set to ‘1’ when the subsystem is not capable of processing NVMe management commands, and the rest of the transmission may be invalid. If cleared to ‘0’, then the NVM Subsystem is fully powered and ready to respond to management commands. This logic level intentionally identifies and prioritizes powered up and ready drives over their powered off neighbors on the same SMBus segment.</p> <p>Drive Functional – Bit 5 is set to ‘1’ to indicate an NVM Subsystem is functional. If cleared to ‘0’, then there is an unrecoverable failure in the NVM Subsystem and the rest of the transmission may be invalid. <u>Note that this bit may default to ‘0’ after reset and transition to ‘1’ after the NVM Subsystem has completed initialization and this case should not be considered an error.</u></p> <p>Reset Not Required - Bit 4 is set to ‘1’ to indicate the NVM Subsystem does not need a reset to resume normal operation. If cleared to ‘0’ then the NVM Subsystem has experienced an error that prevents continued normal operation. A Controller Level Reset is required to resume normal operation.</p> <p>Port 0 PCIe Link Active - Bit 3 is set to ‘1’ to indicate the first port’s PCIe link is up (i.e., the Data Link Control and Management State Machine is in the DL_Active state). If cleared to ‘0’, then the PCIe link is down.</p> <p>Port 1 PCIe Link Active - Bit 2 is set to ‘1’ to indicate the second port’s PCIe link is up. If cleared to ‘0’, then the second port’s PCIe link is down or not present.</p> <p>Bits 1-0 shall be set to ‘1’.</p>

Modify a portion of Section 5.1 (Configuration Get) as follows:

NVMe-MI Configurations are listed in Figure 40. ~~Specifying a reserved identifier in the Configuration Identifier field causes the command to complete with an Invalid Parameter error status.~~

Modify a portion of Figure 17 (Response Message Status Values) as follows:

04h	<p>Invalid Parameter: Invalid command parameter field value. <u>Request Messages received with reserved values in defined fields shall be completed with an Invalid Parameter Error Response Message. Request Messages received with reserved or unimplemented values in defined fields shall be completed with an Invalid Parameter Error Response Message. Other error conditions that result in Invalid Parameter Error Response Message are noted elsewhere in this specification.</u></p>	Refer to 4.2.2
-----	---	----------------

05h	<p>Invalid Command Size: The Command Message body was larger or smaller than that expected by the command due to a reason other than too much or too little input data (e.g., the command did not contain all the required parameters or no input data was expected but the command message body is larger than that needed to contain the required parameters).</p> <p>The expected command message body size is determined by the command opcode assuming no other errors are detected (e.g., Invalid Command Opcode or Invalid ParameterField).</p>	Refer to 4.2.1
-----	---	----------------

Modify a portion of Section 6.2 (Status) as follows:

If the NVMe Admin Command Request Message is well formed, then a success Response Message is transmitted. The success response contains the status associated with NVMe Admin Command in the Status Field of Completion Queue Entry Dword 3. The Status Field contains any NVMe specific status codes (e.g., Success or Invalid ~~ParameterField in Command~~).

Modify a portion of Section 8.2.1 (Controller Metadata) as follows:

If a Set Features command is submitted for this Feature, a Host Metadata data structure, defined in Figure 102, is transferred in the data buffer for the command. The Host Metadata data structure is 4096 bytes in size and contains one or more Metadata Element Descriptors. If host software attempts to add or update a Metadata Element that causes the stored Host Metadata data structure to grow larger than 4096 bytes, the Controller shall abort the command with the status code Invalid ~~ParameterField in Command~~. The Host Metadata structure for this feature is independent of the Host Metadata data structure for the Namespace Metadata feature described in section 8.2.2.

Modify a portion of Section 8.2.2 (Namespace Metadata) as follows:

If a Set Features command is submitted for this Feature, a Host Metadata data structure, defined in Figure 102, is transferred in the data buffer for the command. The Host Metadata data structure is 4096 bytes in size and contains one or more Metadata Element Descriptors. If host software attempts to add or update a Metadata Element that causes the stored Host Metadata data structure to grow larger than 4096 bytes, the Controller shall abort the command with the status code Invalid ~~ParameterField in Command~~. The Host Metadata structure for this feature is independent of the Host Metadata data structure for the Controller Metadata feature described in section 8.2.1.

Modify a portion of Figure 37 (Management Interface Response Message Description) as follows:

Byte	Description
03:00	NVMe-MI Message Header: Refer to 3.2.
04	Status: This field indicates the status of the NVMe-MI command. Refer to 4.2.
07:05	NVMe Management Response: This field is command specific.
N:816	NVMe Response Data (optional)
M+3:M	Message Integrity Check: Refer to 3.2.