



**LEGAL NOTICE:**

© **Copyright 2008 to 2023 NVM Express®, Inc. ALL RIGHTS RESERVED.**

This 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 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: “© 2008 to 2023 NVM Express, Inc. ALL RIGHTS RESERVED.” When making any such citations or references to this document you are not permitted to revise, alter, modify, make any derivatives of, or otherwise amend the referenced portion of this document in any way without the prior express written permission of NVM Express, Inc. Nothing contained in this document shall be deemed as granting you any kind of license to implement or use this document or the specification described therein, or any of its contents, either expressly or impliedly, or to any intellectual property owned or controlled by NVM Express, Inc., including, without limitation, any trademarks of NVM Express, Inc.

**LEGAL DISCLAIMER:**

THIS DOCUMENT AND THE INFORMATION CONTAINED HEREIN IS PROVIDED ON AN “**AS IS**” BASIS. TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, NVM EXPRESS, INC. (ALONG WITH THE CONTRIBUTORS TO THIS DOCUMENT) HEREBY DISCLAIM ALL REPRESENTATIONS, WARRANTIES AND/OR COVENANTS, EITHER EXPRESS OR IMPLIED, STATUTORY OR AT COMMON LAW, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE, VALIDITY, AND/OR NONINFRINGEMENT.

All product names, trademarks, registered trademarks, and/or servicemarks may be claimed as the property of their respective owners.

The NVM Express® design mark is a registered trademark of NVM Express, Inc.

NVM Express Workgroup  
c/o VTM, Inc.  
3855 SW 153<sup>rd</sup> Drive  
Beaverton, OR 97003  
USA  
info@nvmexpress.org

## NVM Express® Technical Proposal (TP)

<b>Technical Proposal ID</b>	TP6021 Status Reporting Enhancements
<b>Revision Date</b>	2023.06.11
<b>Builds on Specification(s)</b>	NVMe Management Interface Specification, Revision 1.2b NVM Express Base Specification, Revision 2.0b
<b>References</b>	TP6027 Reset Behavior Clarifications TP6033 MCTP Packet Timing

### Technical Proposal Author(s)

<b>Name</b>	<b>Company</b>
Austin Bolen	Dell EMC
Myron Loewen	Solidigm

### Technical Proposal Overview

This technical proposal makes a variety of enhancements and clarifications related to NVMe-MI status reporting.

### Revision History

<b>Revision Date</b>	<b>Author</b>	<b>Change Description</b>
2022.01.19	Austin Bolen	<ul style="list-style-type: none"><li>Initial draft.</li></ul>

*Technical input submitted to the NVM Express® Workgroup is subject to the terms of the NVM Express® Participant's agreement. Copyright © 2008 to 2023 NVM Express, Inc.*

Revision Date	Author	Change Description
2022.01.31	Austin Bolen	<ul style="list-style-type: none"> <li>• Updates from workgroup discussion on 1/31.</li> </ul>
2022.02.14	Austin Bolen	<ul style="list-style-type: none"> <li>• Updates from workgroup discussion on 2/14.</li> </ul>
2022.03.12	Austin Bolen	<ul style="list-style-type: none"> <li>• Added most of the content to report out-of-band Telemetry Controller-Initiated Data Available.</li> <li>• Aligned the behavior of the out-of-band NSSRO bits to the CSTS.NSSRO bit.</li> <li>• Editorial updates.</li> </ul>
2022.04.04	Austin Bolen	<ul style="list-style-type: none"> <li>• Wordsmithing.</li> <li>• Added security status changed framework but will remove that and create a separate TPAR.</li> </ul>
2022.04.11	Austin Bolen	<ul style="list-style-type: none"> <li>• Removed the following which are moving to new TPARs: <ul style="list-style-type: none"> <li>○ Security status change reporting.</li> <li>○ Error reporting bits.</li> <li>○ Namespace create/delete status bit.</li> </ul> </li> </ul>
2022.05.23	Austin Bolen	<ul style="list-style-type: none"> <li>• Modified Composite Temperature to allow reporting coldest temp in under-temperature conditions.</li> <li>• Clarified that the critical warning bits in Controller Health Data Structure match the Critical Warning bits in the SMART / Health Information log.</li> <li>• Added technical proposal overview.</li> <li>• Added description for the changes documents.</li> <li>• Editorial updates.</li> </ul>
2022.05.24	Austin Bolen	<ul style="list-style-type: none"> <li>• Added capability bit advertise support for the backwards incompatible changes to the status bits.</li> <li>• Editorial updates.</li> </ul>
2022.06.02	Austin Bolen	<ul style="list-style-type: none"> <li>• Changes to NSSRO bits to remove dependency on CSTS.NSSRO as that bit is optional and may not be implemented even though NSSRO is reported out-of-band.</li> <li>• Updates from TWG meeting. <ul style="list-style-type: none"> <li>○ Wordsmithed Telemetry Controller-Initiated Data Available definition.</li> <li>○ Added note to editor to obsolete the Status Reporting Enhancements in the NVMe-MI 1.2_Next integration file.</li> </ul> </li> </ul>
2022.06.06	Austin Bolen	<ul style="list-style-type: none"> <li>• Added a field in each Telemetry log to indicate if it is Controller in scope or NVM Subsystem in scope.</li> <li>• Added a new error status code to return if host attempts to read the Telemetry Controller-Initiated log when the Telemetry Controller-Initiated Data Available field is 0h.</li> <li>• Only set the status bits for Telemetry Controller-Initiated log in the data structures for each Controller if the log is Controller in scope.</li> </ul>
2022.06.19	Austin Bolen	<ul style="list-style-type: none"> <li>• Editorial updates.</li> <li>• Removed Telemetry Controller-Initiated Data Unavailable error status code.</li> </ul>
2022.06.20	Austin Bolen	<ul style="list-style-type: none"> <li>• Editorial updates.</li> <li>• Removed Telemetry Controller-Initiated Data Unavailable error status code from the new requirements in the changes document.</li> <li>• Specified that a Controller Health Status Poll command with the Clear Changed Flags bit set to '1' shall clear the Namespace Attribute Changed and Firmware Activated bits in Controller Health Data Structure to '0'.</li> </ul>
2022.07.11	Austin Bolen	<ul style="list-style-type: none"> <li>• Addressed phase 3 comments.</li> </ul>
2022.07.21	Austin Bolen	<ul style="list-style-type: none"> <li>• Don't obsolete in the Status Reporting Enhancements bit in the next version of the NVMe-MI specification.</li> </ul>
2022.08.29	Austin Bolen	<ul style="list-style-type: none"> <li>• Editorial updates based on member review feedback and workgroup discussions.</li> </ul>
2022.09.02	Austin Bolen	<ul style="list-style-type: none"> <li>• Fixed issue from last update.</li> </ul>
2022.12.06	Devin Allison	<ul style="list-style-type: none"> <li>• Integrated</li> </ul>
2023.01.15	Devin Allison	<ul style="list-style-type: none"> <li>• Integrated comments from Austin Bolen, Myron Loewen, and Mike Allison review</li> </ul>
2023.01.22	Devin Allison	<ul style="list-style-type: none"> <li>• Integrated comments from Austin Bolen and Mike Allison review</li> </ul>
2023.01.29	Devin Allison	<ul style="list-style-type: none"> <li>• Integrated comments from Austin Bolen and Mike Allison after MI review</li> </ul>
2023.02.07	Devin Allison	<ul style="list-style-type: none"> <li>• Integrated comments from Austin Bolen and Mike Allison after MI review</li> </ul>

Revision Date	Author	Change Description
2023.02.27	Austin Bolen	<ul style="list-style-type: none"> <li>• Changed a few instances of “should” to “shall”.</li> <li>• Added a requirement that the steps in Clear Changed Flags be done atomically and in the order listed.</li> <li>• Fixed how the Telemetry Controller-Initiated Data Available status bits work in the out-of-band mechanism.</li> <li>• Fixed a “field” that should have been a “bit”.</li> </ul>
2023.06.06	Devin Allison	<ul style="list-style-type: none"> <li>• Integrated Austin Bolen and Mike Allison comments</li> </ul>

## Description for Changes Document for NVM Express Management Interface Specification

### Feature Enhancements:

- Updates to Telemetry logs (Mandatory if Telemetry logs are supported)
  - **New Requirement:** Added a mechanism to report Telemetry Controller-Initiated Data Available via the out-of-band mechanism.
  - **New Requirement:** Added a field in each Telemetry log to indicate if it is Controller in scope or NVM Subsystem in scope.
- **New Requirement:** Added a status bit to indicate the NVM Subsystem is in the Sanitize failure mode (Mandatory if Sanitize is supported)
- Updates to Controller Health Status Poll (Mandatory)
  - **New Requirement:** Added a requirement to return the Controller Health Status Changed Flags in the response.
  - Clarified how the Controller filtering logic works and reversed the polarity to indicate when Controllers meet the selection criteria to be included in the Controller Health Status Poll response versus when the Controllers are filtered out.
  - Clarified that it is not an error if the INCF bit is set to ‘1’ and non-SR-IOV PCI Functions do not exist.
  - Clarified that if the INCF, INCPF, or INCVF bits are cleared to ‘0’, then the corresponding Controller is not included in the Controller Health Status Poll response.
  - **New Requirement:** Specified that a Controller ID that is greater than the maximum Controller ID in the NVM Subsystem results in an Invalid Parameter Error Response.
  - Clarified that if the Clear Error State Flags bit in Get State is set to ‘1’, then the bits are cleared after, not before, the current state of the bits are copied to the response and the Controller selection criteria based on Controller Health Status Changed flags have been performed.
  - **Incompatible Change:** Changed the Controller Enable Change Occurred bit in the Controller Health Data Structure from a status bit that is set when the Controller is enabled or disabled to a state bit that indicates if the Controller is currently enabled or disabled.
- Updates to the Controller Health Status Changed Flags (Mandatory)
  - **Incompatible Change:** Set the Controller Enable Change Occurred bit whenever Controller Enable Change Occurred bit in the Controller Health Data Structure changes state instead of only when it transitions from ‘0’ to ‘1’.
  - **Incompatible Change:** Set the Critical Warning bit whenever any bits in the Critical Warning field change state instead of only when they transition from ‘0’ to ‘1’.
  - **Incompatible Change:** Set the Controller Status Change bit whenever any bits in the Controller Status field change state instead of only when they transition from ‘0’ to ‘1’.
  - **Incompatible Change:** Set the Controller Fatal Status bit whenever Controller Fatal Status bit in the Controller Health Data Structure changes state instead of only when it transitions from ‘0’ to ‘1’.
  - **Incompatible Change:** Set the Ready bit whenever Ready bit in the Controller Health Data Structure changes state instead of only when it transitions from ‘0’ to ‘1’.
  - **Incompatible Change:** Specified that a Controller Health Status Poll command with the Clear Changed Flags bit set to ‘1’ shall also clear the Namespace Attribute Changed and Firmware Activated bits in Controller Health Data Structure to ‘0’.

- Updates to Get State (Mandatory)
  - Clarified that if the Clear Changed Flags bit is set to '1, then the bits are cleared after, not before, the current state of the bits are copied to the response.
  - **New Requirement:** Specified that NVM Subsystem Reset Occurred is not cleared by a Management Endpoint Reset.
- Updates to NVM Subsystem Health Status Poll (Mandatory)
  - **New Requirement:** Added the ability report the Composite Temperature of the coldest Controller in the NVM Subsystem for under-temperature conditions.
  - Clarified the condition that causes a temperature sensor failure to be reported (if the composite temperature of the NVM Subsystem cannot be guaranteed to be accurate due to the failure of one or more temperature sensors).
  - Clarified that if the Clear Status bit is set to '1, then the bits are cleared after, not before, the current state of the bits are copied to the response.
- NVM Subsystem Reset Occurred updates (Mandatory)
  - **Incompatible Change:** Specified in Get State that the NVM Subsystem Reset Occurred status bits are only set if a new firmware image is not activated to synchronize to the behavior of NVM Subsystem Reset Occurred in the NVM Express Base Specification.

## Description for Changes Document for NVM Express Base Specification

### Feature Enhancements:

- Added a mechanism to report Telemetry Controller-Initiated Data Available via the out-of-band mechanism (Mandatory if Telemetry logs are supported)

### Markup Conventions:

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

## Description of Specification Changes for NVM Express Management Interface Specification

Modify section 4.2.1.4 as follows:

### 4.2.1.4 Get State

The Get State Control Primitive is used to get and clear the state of a ~~Command Slot and~~ Management Endpoint and the Management Endpoint's Command Slots.

The Management Endpoint shall contain Management Endpoint state as shown in Figure **TBD1**.

Figure **TBD1**: Management Endpoint State Data Structure

Bytes	Description		
1:0	<b>Management Endpoint State (MES):</b> This field shall indicate the Management Endpoint state. A Management Endpoint Reset of the corresponding Management Endpoint shall clear bit 15 to '0' and clear bits 13:0 to 0h.		
	Bits	Command Slot Specific <sup>1</sup>	Description
	15	Yes	<b>Pause Flag (PFLG):</b> This bit indicates whether or not the Command Slot is paused. This bit set to '1' indicates the Command Slot is paused. This bit cleared to '0' indicates the Command Slot is not paused.  While the Pause Flag is set, the Management Endpoint disables the timeout waiting for packet timer (refer to section 4.2.1.1) for the Command Slot and does not transmit responses to Command Messages.
	14	No	<b>NVM Subsystem Reset Occurred (NSSRO):</b> <del>This bit indicates when</del> If an NVM Subsystem Reset occurs due to any reason other than application of main power and does not cause activation of a new firmware image, then this <del>bit shall be set to '1' if the last occurrence of an NVM Subsystem Reset occurred while main power was applied to the NVM Subsystem.</del>  If an NVM Subsystem Reset occurs due to application of main power or causes activation of a new firmware image, then this bit shall be cleared to '0'. <del>This bit is cleared to '0' following a power cycle and following a Get State Control Primitive with the CESF bit set to '1'.</del>
	13	No	<b>Bad Packet or Other Physical Layer (BPOPL):</b> If the Management Endpoint detects an error of this type (refer to the MCTP Base Specification), then this bit shall be set to '1'. <del>This bit is set to '1' if a packet sent to the Management Endpoint failed a transport specific packet integrity check since the last time Get State Control Primitive was processed with the CESF bit set to '1'.</del>
	12	No	<b>Bad, Unexpected, or Expired Message Tag (BUEMT):</b> <del>This bit is set to '1' if the Management Endpoint detects</del> ed an error of this type (refer to the MCTP Base Specification), then this bit shall be set to '1' <del>since the last time Get State Control Primitive was processed with the CESF bit set to '1'.</del>
	11	No	<b>Out-of-Sequence Packet Sequence Number (OSPSN):</b> <del>This bit is set to '1' if the Management Endpoint detects</del> ed an error of this type (refer to the MCTP Base Specification), then this bit shall be set to '1' <del>since the last time Get State Control Primitive was processed with the CESF bit set to '1'.</del>
	10	No	<b>Unexpected Middle or End of Packet (UMEP):</b> <del>This bit is set to '1' if the Management Endpoint detects</del> ed an error of this type (refer to the MCTP Base Specification), then this bit shall be set to '1' <del>since the last time Get State Control Primitive was processed with the CESF bit set to '1'.</del>

Figure TBD1: Management Endpoint State Data Structure

Bytes		Description										
09	No	<b>Incorrect Transmission Unit (ITU):</b> <del>This bit is set to '1' if</del> the Management Endpoint detects <del>ed</del> an error of this type (refer to the MCTP Base Specification), then this bit shall be set to '1' <del>since the last time Get State Control Primitive was processed with the CESF bit set to '1'.</del>										
08	No	<b>Unknown Destination ID (UDSTID):</b> <del>This bit is set to '1' if</del> the Management Endpoint detects <del>ed</del> an error of this type (refer to the MCTP Base Specification), then this bit shall be set to '1' <del>since the last time Get State Control Primitive was processed with the CESF bit set to '1'.</del>										
07	No	<b>Bad Header Version (BHVS):</b> <del>This bit is set to '1' if</del> the Management Endpoint detects <del>ed</del> an error of this type (refer to the MCTP Base Specification), then this bit shall be set to '1' <del>since the last time Get State Control Primitive was processed with the CESF bit set to '1'.</del>										
06	No	<b>Unsupported Transmission Unit (UTUNT):</b> <del>This bit is set to '1' if</del> the Management Endpoint detects <del>ed</del> an error of this type (refer to the MCTP Base Specification), then this bit shall be set to '1' <del>since the last time Get State Control Primitive was processed with the CESF bit set to '1'.</del>										
...	...	...										
04	No	<b>Bad Message Integrity Check Error (BMICE):</b> <del>This bit is set to '1' if</del> the Management Endpoint detects <del>ed</del> an error of this type (refer to the MCTP Base Specification), then this bit shall be set to '1' <del>since the last time Get State Control Primitive was processed with the CESF bit set to '1'.</del>										
03	No	<b>Command Message to non-Idle Command Slot (CMNICS):</b> <del>This bit is set to '1'</del> if the Management Endpoint discards <del>ed</del> one or more Command Messages due to overlapping Command Messages to a Command Slot, then this bit shall be set to '1' <del>since the last time Get State Control Primitive was processed with the CESF bit set to '1'.</del>										
02		Reserved										
01:00	Yes	<b>Slot Command Servicing State (SSTA):</b> This field shall indicate <del>s</del> the current command servicing state of the Command Slot. An implementation may choose to indicate only the Idle and Process states in this field. Refer to <a href="#">Figure 33</a> . <table border="1" data-bbox="824 1171 1149 1318"> <thead> <tr> <th>Value</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0h</td> <td>Idle</td> </tr> <tr> <td>1h</td> <td>Receive</td> </tr> <tr> <td>2h</td> <td>Process</td> </tr> <tr> <td>3h</td> <td>Transmit</td> </tr> </tbody> </table>	Value	Description	0h	Idle	1h	Receive	2h	Process	3h	Transmit
Value	Description											
0h	Idle											
1h	Receive											
2h	Process											
3h	Transmit											
<b>NOTES:</b> 1. Command Slot Specific. A 'Yes' in this column indicates the value of the field is independent per Command Slot within a Management Endpoint. A 'No' in this column indicates the same value is reported for either Command Slot.												

The format of the CPSP field in the Control Primitive Request Message is shown in Figure 41.

**Figure 41: Get State Control Primitive Request Message Fields**

Bytes	Description					
07:06	<b>Control Primitive Specific Parameter (CPSP):</b> This field specifies <del>is used to pass</del> Control Primitive specific parameter information.					
	<table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>15:01</td> <td>Reserved</td> </tr> <tr> <td>00</td> <td> <b>Clear Error State Flags (CESF):</b> If this bit is set to '1', then the Management Endpoint shall perform the following steps atomically in the order listed:                             <ol style="list-style-type: none"> <li>copy the current value of the Management Endpoint State data structure (refer to <a href="#">Figure TBD1</a>) to the Control Primitive Specific Response field of the Response Message (refer to <a href="#">Figure 42</a>); and</li> <li>clear bits 14:03 in the Management Endpoint State data structure to 0h<del>the error state flags</del>.</li> </ol>                               If this bit is cleared to '0', then the Management Endpoint shall copy the current value of the Management Endpoint State data structure (refer to <a href="#">Figure TBD1</a>) to the Control Primitive Specific Response field of the Response Message (refer to <a href="#">Figure 42</a>) and shall not clear modify bits 14:03 in the Management Endpoint State data structure<del>the error state flags</del>.                             </td> </tr> </tbody> </table>	Bits	Description	15:01	Reserved	00
Bits	Description					
15:01	Reserved					
00	<b>Clear Error State Flags (CESF):</b> If this bit is set to '1', then the Management Endpoint shall perform the following steps atomically in the order listed: <ol style="list-style-type: none"> <li>copy the current value of the Management Endpoint State data structure (refer to <a href="#">Figure TBD1</a>) to the Control Primitive Specific Response field of the Response Message (refer to <a href="#">Figure 42</a>); and</li> <li>clear bits 14:03 in the Management Endpoint State data structure to 0h<del>the error state flags</del>.</li> </ol> If this bit is cleared to '0', then the Management Endpoint shall copy the current value of the Management Endpoint State data structure (refer to <a href="#">Figure TBD1</a> ) to the Control Primitive Specific Response field of the Response Message (refer to <a href="#">Figure 42</a> ) and shall not clear modify bits 14:03 in the Management Endpoint State data structure <del>the error state flags</del> .					

~~The Management Endpoint shall transmit a Response Message with success status after receiving the Get State Control Primitive.~~ The format of the CPSP field in the Control Primitive Success Response is shown in Figure 42.

**Figure 42: Get State Control Primitive Success Response Fields**

Bytes	Description																					
07:06	<b>Control Primitive Specific Response (CPSR):</b> This field is used to return Control Primitive specific status.																					
	<table border="1"> <thead> <tr> <th>Bits</th> <th>Command Slot Specific<sup>4</sup></th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>45</td> <td>Yes</td> <td> <del><b>Pause Flag (PFLG):</b> This bit indicates whether or not the Command Slot is paused. This bit set to '1' indicates the Command Slot is paused. This bit cleared to '0' indicates the Command Slot is not paused. While the Pause Flag is set, the Management Endpoint disables the timeout waiting for packet timer (refer to section 4.2.1.1) for the Command Slot and does not transmit responses to Command Messages.</del> </td> </tr> <tr> <td>44</td> <td>No</td> <td> <del><b>NVM Subsystem Reset Occurred (NSSRO):</b> This bit indicates when an NVM Subsystem Reset occurs while main power is applied. This bit is set to '1' if the last occurrence of an NVM Subsystem Reset occurred while main power was applied to the NVM Subsystem. This bit is cleared to '0' following a power cycle and following a Get State Control Primitive with the CESF bit set to '1'.</del> </td> </tr> <tr> <td>43</td> <td>No</td> <td> <del><b>Bad Packet or Other Physical Layer (BPOPL):</b> This bit is set to '1' if a packet sent to the Management Endpoint failed a transport specific packet integrity check since the last time Get State Control Primitive was processed with the CESF bit set to '1'.</del> </td> </tr> <tr> <td>42</td> <td>No</td> <td> <del><b>Bad, Unexpected, or Expired Message Tag (BUEMT):</b> This bit is set to '1' if the Management Endpoint detected an error of this type (refer to the MCTP Base Specification) since the last time Get State Control Primitive was processed with the CESF bit set to '1'.</del> </td> </tr> <tr> <td>41</td> <td>No</td> <td> <del><b>Out-of-Sequence Packet Sequence Number (OSPSN):</b> This bit is set to '1' if the Management Endpoint detected an error of this type (refer to the MCTP Base Specification) since the last time Get State Control Primitive was processed with the CESF bit set to '1'.</del> </td> </tr> <tr> <td>40</td> <td>No</td> <td> <del><b>Unexpected Middle or End of Packet (UMEP):</b> This bit is set to '1' if the Management Endpoint detected an error of this type (refer to the MCTP Base Specification) since the last time Get State Control Primitive was processed with the CESF bit set to '1'.</del> </td> </tr> </tbody> </table>	Bits	Command Slot Specific <sup>4</sup>	Description	45	Yes	<del><b>Pause Flag (PFLG):</b> This bit indicates whether or not the Command Slot is paused. This bit set to '1' indicates the Command Slot is paused. This bit cleared to '0' indicates the Command Slot is not paused. While the Pause Flag is set, the Management Endpoint disables the timeout waiting for packet timer (refer to section 4.2.1.1) for the Command Slot and does not transmit responses to Command Messages.</del>	44	No	<del><b>NVM Subsystem Reset Occurred (NSSRO):</b> This bit indicates when an NVM Subsystem Reset occurs while main power is applied. This bit is set to '1' if the last occurrence of an NVM Subsystem Reset occurred while main power was applied to the NVM Subsystem. This bit is cleared to '0' following a power cycle and following a Get State Control Primitive with the CESF bit set to '1'.</del>	43	No	<del><b>Bad Packet or Other Physical Layer (BPOPL):</b> This bit is set to '1' if a packet sent to the Management Endpoint failed a transport specific packet integrity check since the last time Get State Control Primitive was processed with the CESF bit set to '1'.</del>	42	No	<del><b>Bad, Unexpected, or Expired Message Tag (BUEMT):</b> This bit is set to '1' if the Management Endpoint detected an error of this type (refer to the MCTP Base Specification) since the last time Get State Control Primitive was processed with the CESF bit set to '1'.</del>	41	No	<del><b>Out-of-Sequence Packet Sequence Number (OSPSN):</b> This bit is set to '1' if the Management Endpoint detected an error of this type (refer to the MCTP Base Specification) since the last time Get State Control Primitive was processed with the CESF bit set to '1'.</del>	40	No	<del><b>Unexpected Middle or End of Packet (UMEP):</b> This bit is set to '1' if the Management Endpoint detected an error of this type (refer to the MCTP Base Specification) since the last time Get State Control Primitive was processed with the CESF bit set to '1'.</del>
	Bits	Command Slot Specific <sup>4</sup>	Description																			
	45	Yes	<del><b>Pause Flag (PFLG):</b> This bit indicates whether or not the Command Slot is paused. This bit set to '1' indicates the Command Slot is paused. This bit cleared to '0' indicates the Command Slot is not paused. While the Pause Flag is set, the Management Endpoint disables the timeout waiting for packet timer (refer to section 4.2.1.1) for the Command Slot and does not transmit responses to Command Messages.</del>																			
	44	No	<del><b>NVM Subsystem Reset Occurred (NSSRO):</b> This bit indicates when an NVM Subsystem Reset occurs while main power is applied. This bit is set to '1' if the last occurrence of an NVM Subsystem Reset occurred while main power was applied to the NVM Subsystem. This bit is cleared to '0' following a power cycle and following a Get State Control Primitive with the CESF bit set to '1'.</del>																			
	43	No	<del><b>Bad Packet or Other Physical Layer (BPOPL):</b> This bit is set to '1' if a packet sent to the Management Endpoint failed a transport specific packet integrity check since the last time Get State Control Primitive was processed with the CESF bit set to '1'.</del>																			
	42	No	<del><b>Bad, Unexpected, or Expired Message Tag (BUEMT):</b> This bit is set to '1' if the Management Endpoint detected an error of this type (refer to the MCTP Base Specification) since the last time Get State Control Primitive was processed with the CESF bit set to '1'.</del>																			
	41	No	<del><b>Out-of-Sequence Packet Sequence Number (OSPSN):</b> This bit is set to '1' if the Management Endpoint detected an error of this type (refer to the MCTP Base Specification) since the last time Get State Control Primitive was processed with the CESF bit set to '1'.</del>																			
40	No	<del><b>Unexpected Middle or End of Packet (UMEP):</b> This bit is set to '1' if the Management Endpoint detected an error of this type (refer to the MCTP Base Specification) since the last time Get State Control Primitive was processed with the CESF bit set to '1'.</del>																				

**Figure 42: Get State Control Primitive Success Response Fields**

Bytes	Description												
	09	No	<del><b>Incorrect Transmission Unit (ITU):</b> This bit is set to '1' if the Management Endpoint detected an error of this type (refer to the MCTP Base Specification) since the last time Get State Control Primitive was processed with the CESF bit set to '1'.</del>										
	08	No	<del><b>Unknown Destination ID (UDSTID):</b> This bit is set to '1' if the Management Endpoint detected an error of this type (refer to the MCTP Base Specification) since the last time Get State Control Primitive was processed with the CESF bit set to '1'.</del>										
	07	No	<del><b>Bad Header Version (BHVS):</b> This bit is set to '1' if the Management Endpoint detected an error of this type (refer to the MCTP Base Specification) since the last time Get State Control Primitive was processed with the CESF bit set to '1'.</del>										
	06	No	<del><b>Unsupported Transmission Unit (UTUNT):</b> This bit is set to '1' if the Management Endpoint detected an error of this type (refer to the MCTP Base Specification) since the last time Get State Control Primitive was processed with the CESF bit set to '1'.</del>										
	05	No	<del><b>Timeout Waiting for a Packet (WPTT):</b> This bit is set to '1' if the Management Endpoint detected an error of this type (refer to the MCTP Base Specification) since the last time Get State Control Primitive was processed with the CESF bit set to '1'.</del>										
	04	No	<del><b>Bad Message Integrity Check Error (BMICE):</b> This bit is set to '1' if the Management Endpoint detected an error of this type (refer to the MCTP Base Specification) since the last time Get State Control Primitive was processed with the CESF bit set to '1'.</del>										
	03	No	<del><b>Command Message to non-Idle Command Slot (CMNICS):</b> This bit is set to '1' if the Management Endpoint discarded one or more Command Messages due to overlapping Command Messages to a Command Slot since the last time Get State Control Primitive was processed with the CESF bit set to '1'.</del>										
	02		Reserved										
	01:00	Yes	<del><b>Slot Command Servicing State (SSTA):</b> This field indicates the current command servicing state of the Command Slot. An implementation may choose to indicate only the Idle and Process states in this field. Refer to Figure 33.</del> <table border="1" data-bbox="820 1081 1144 1228"> <thead> <tr> <th>Value</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0h</td> <td>Idle</td> </tr> <tr> <td>1h</td> <td>Receive</td> </tr> <tr> <td>2h</td> <td>Process</td> </tr> <tr> <td>3h</td> <td>Transmit</td> </tr> </tbody> </table>	Value	Description	0h	Idle	1h	Receive	2h	Process	3h	Transmit
Value	Description												
0h	Idle												
1h	Receive												
2h	Process												
3h	Transmit												
			<b>NOTES:</b> 1. <del>Command Slot Specific. A 'Yes' in this column indicates the value of the field is independent per Command Slot within a Management Endpoint. A 'No' in this column indicates the same value is reported for either Command Slot.</del>										

**Modify Section 5.1.2 as follows:**

### 5.1.2 Health Status Change (Configuration Identifier 02h)

The Health Status Change configuration is used to clear the selected status bits in the Composite Controller Status [Flags](#) field using Configuration Set. A Requester should not use Configuration Get for this Configuration Identifier.

The configuration specific fields in [the NVMe Management Dwords 0 and 1 fields](#) are reserved. A Responder shall complete a Configuration Get command on this Configuration Identifier with a Success Response. The NVMe Management Response field is reserved and there is no Response Data.

**Modify Section 5.2.2 as follows:**

## 5.2.2 Health Status Change (Configuration Identifier 02h)

This Configuration Identifier is used to clear selected status bits in the Composite Controller Status **Flags** field ~~of the NVM Subsystem Health Data Structure (refer to Figure 90) returned by the NVM Subsystem Health Status Poll command.~~

The Composite Controller Status **Flags** field ~~of the NVM Subsystem Health Data Structure~~ is used to report the occurrence of health and status events associated with the NVM Subsystem ~~via the Composite Controller Status field in the Response Message for the NVM Subsystem Health Status Poll command.~~ When a bit in this field is set to '1', ~~it that bit remains a set to '1' until cleared to '0' by a Requester or until cleared to '0' by a reset as described in Figure TBD2.~~

A Configuration Set command that selects Health Status Change ~~may be used to clear~~ clears corresponding bits selected in NVMe Management Dword 1 of the Composite Controller Status **Flags** field to '0'.

A Configuration Set command that selects Health Status Change operates independently in the out-of-band mechanism and the in-band tunneling mechanism.

An NVMe Storage Device or NVMe Enclosure supporting the Health Status Change Configuration Identifier in the out-of-band mechanism shall have an independent copy of the Composite Controller Status **Flags** dedicated to the out-of-band mechanism. In the out-of-band mechanism, a Configuration Set command that selects Health Status Change only applies to the copy of the Composite Controller Status **Flags** dedicated to the out-of-band mechanism. Refer to section 5.4 for more details on Composite Controller Status **Flags**.

An NVMe Storage Device or NVMe Enclosure supporting the Health Status Change Configuration Identifier in the in-band tunneling mechanism shall have an independent copy of the Composite Controller Status **Flags** dedicated to the in-band tunneling mechanism. In the in-band tunneling mechanism, a Configuration Set command that selects Health Status Change only applies to the copy of the Composite Controller Status **Flags** dedicated to the in-band tunneling mechanism.

**Figure 73: Health Status Change - NVMe Management Dword 0**

Bits	Description
31:08	Reserved
07:00	<b>Configuration Identifier:</b> This field specifies the identifier of the Configuration that is being written. Refer to Figure 65.

**Figure 74: Health Status Change – NVMe Management Dword 1**

Bits	Description
31: <del>13</del> <del>42</del>	Reserved
12	<b>Telemetry Controller-Initiated Data Available (TCIDA):</b> If this bit is set to '1', then bit 13 in the Composite Controller Status <b>Flags</b> field shall be cleared to '0'. If this bit is cleared to '0', then bit 13 in the Composite Controller Status <b>Flags</b> field shall not be modified.
11	<b>Critical Warning (CWARN):</b> <del>When</del> If this bit is set to '1', then bit 12 in the Composite Controller Status <b>Flags</b> field <del>of the NVM Subsystem Health Data Structure is</del> shall be cleared to '0'. If this bit is cleared to '0', then bit 12 in the Composite Controller Status <b>Flags</b> field shall not be modified.
10	<b>Available Spare (SPARE):</b> <del>When</del> If this bit is set to '1', then bit 11 in the Composite Controller Status <b>Flags</b> field <del>of the NVM Subsystem Health Data Structure is</del> shall be cleared to '0'. If this bit is cleared to '0', then bit 11 in the Composite Controller Status <b>Flags</b> field shall not be modified.

**Figure 74: Health Status Change – NVMe Management Dword 1**

Bits	Description
09	<b>Percentage Used (PDLU):</b> <del>When</del> If this bit is set to '1', then bit 10 in the Composite Controller Status Flags field <del>of the NVMe Subsystem Health Data Structure is</del> shall be cleared to '0'. If this bit is cleared to '0', then bit 10 in the Composite Controller Status Flags field shall not be modified.
08	<b>Composite Temperature (CTEMP):</b> <del>When</del> If this bit is set to '1', then bit 9 in the Composite Controller Status Flags field <del>of the NVMe Subsystem Health Data Structure is</del> shall be cleared to '0'. If this bit is cleared to '0', then bit 9 in the Composite Controller Status Flags field shall not be modified.
07	<b>Controller Status Change (CSCHNG):</b> <del>When</del> If this bit is set to '1', then bit 8 in the Composite Controller Status Flags field <del>of the NVMe Subsystem Health Data Structure is</del> shall be cleared to '0'. If this bit is cleared to '0', then bit 8 in the Composite Controller Status Flags field shall not be modified.
06	<b>Firmware Activated (FA):</b> <del>When</del> If this bit is set to '1', then bit 7 in the Composite Controller Status Flags field <del>of the NVMe Subsystem Health Data Structure is</del> shall be cleared to '0'. If this bit is cleared to '0', then bit 7 in the Composite Controller Status Flags field shall not be modified.
05	<b>Namespace Attribute Changed (NAC):</b> <del>When</del> If this bit is set to '1', then bit 6 in the Composite Controller Status Flags field <del>of the NVMe Subsystem Health Data Structure is</del> shall be cleared to '0'. If this bit is cleared to '0', then bit 6 in the Composite Controller Status Flags field shall not be modified.
04	<b>Controller Enable Change Occurred (CECO):</b> <del>When</del> If this bit is set to '1', then bit 5 in the Composite Controller Status Flags field <del>of the NVMe Subsystem Health Data Structure is</del> shall be cleared to '0'. If this bit is cleared to '0', then bit 5 in the Composite Controller Status Flags field shall not be modified.
03	<b>NVM Subsystem Reset Occurred (NSSRO):</b> <del>When</del> If this bit is set to '1', then bit 4 in the Composite Controller Status Flags field <del>of the NVMe Subsystem Health Data Structure is</del> shall be cleared to '0'. If this bit is cleared to '0', then bit 4 in the Composite Controller Status Flags field shall not be modified.
02	<b>Shutdown Status (SHST):</b> <del>When</del> If this bit is set to '1', then bit 2 in the Composite Controller Status Flags field <del>of the NVMe Subsystem Health Data Structure is</del> shall be cleared to '0'. If this bit is cleared to '0', then bit 2 in the Composite Controller Status Flags field shall not be modified.
01	<b>Controller Fatal Status (CFS):</b> <del>When</del> If this bit is set to '1', then bit 1 in the Composite Controller Status Flags field <del>of the NVMe Subsystem Health Data Structure is</del> shall be cleared to '0'. If this bit is cleared to '0', then bit 1 in the Composite Controller Status Flags field shall not be modified.
00	<b>Ready (RDY):</b> <del>When</del> If this bit is set to '1', then bit 0 in the Composite Controller Status Flags field <del>of the NVMe Subsystem Health Data Structure is</del> shall be cleared to '0'. If this bit is cleared to '0', then bit 0 in the Composite Controller Status Flags field shall not be modified.

**Modify Section 5.3 as follows:**

### 5.3 Controller Health Status Poll

The Controller Health Status Poll command is used to efficiently determine changes in health status attributes associated with one or more Controllers in the NVM Subsystem. [This command returns a list of zero or more Controller Health Data Structures based on various selection criteria \(refer to section 5.3.1\).](#)

The Controller Health Status Poll command operates independently in the out-of-band mechanism and the in-band tunneling mechanism.

An NVMe Storage Device or NVMe Enclosure supporting the Controller Health Status Poll command in the out-of-band mechanism shall have an independent copy of the Controller Health Data Structure (refer to Figure 80) and the Controller Health Status Changed Flags (refer to Figure 81) dedicated to the out-of-band mechanism. In the out-of-band mechanism, a Controller Health Status Poll command only applies to the copy of the Controller Health Data Structure and the Controller Health Status Changed Flags dedicated to the out-of-band mechanism.

An NVMe Storage Device or NVMe Enclosure supporting the Controller Health Status Poll command in the in-band tunneling mechanism shall have an independent copy of the Controller Health Data Structure and the Controller Health Status Changed Flags dedicated to the in-band tunneling mechanism. In the in-band tunneling mechanism, a Controller Health Status Poll command only applies to the copy of the Controller Health Data Structure and the Controller Health Status Changed Flags dedicated to the in-band tunneling mechanism.

The Controller Health Status Poll command uses the NVMe Management Dwords 0 and 1 fields. The format of the NVMe Management Dword 0 field is shown in Figure 77 and the format of the NVMe Management Dword 1 field is shown in Figure 78.

**Figure 77: Controller Health Status Poll – NVMe Management Dword 0**

Bits	Description
31	<p><b>Report All (ALL):</b> <del>When-If</del> this bit is set to '1', <del>then a Controller Health Data Structure is returned regardless of the status of the Controller Health Status Changed Flags. The Controller selection fields (SCTLID, MAXRENT, INCF, INCPF, and INCVF) still apply even when this bit is set to '1' but</del> the error selection bits (CWARN, SPARE, PDLU, CTEMP, and CSTS in Figure 78) <del>shall be ignored when determining whether to return the Controller Health Data Structure per the selection criteria in section 5.3.1 do not apply.</del></p> <p><del>When-If</del> this bit is cleared to '0', <del>then the a Controller Health Data Structure is returned based on the Controller selection fields (SCTLID, MAXRENT, INCF, INCPF, and INCVF) and</del> error selection bits fields (CWARN, SPARE, PDLU, CTEMP, and CSTS in Figure 78) <del>shall not be ignored when determining whether to return the Controller Health Data Structure per the selection criteria in section 5.3.1.</del></p>
30:27	Reserved
26	<p><b>Include SR-IOV Virtual Functions (INCVF):</b> <del>When-If</del> this bit is set to '1', then a Controller Health Data Structure <del>is shall be</del> returned for Controllers associated with SR-IOV Virtual Functions (VFs) <del>unless excluded by other selection criteria as described in section 5.3.1.</del></p> <p>If this bit is cleared to '0', then a Controller Health Data Structure shall not be returned for Controllers associated with SR-IOV VFs.</p> <p>It is not an error if this bit is set to '1' and SR-IOV Virtual <del>f</del>Functions do not exist.</p>
25	<p><b>Include SR-IOV Physical Functions (INCPF):</b> <del>When-If</del> this bit is set to '1', then a Controller Health Data Structure <del>is shall be</del> returned for Controllers associated with SR-IOV Physical Functions (PFs) <del>unless excluded by other selection criteria as described in section 5.3.1.</del></p> <p>If this bit is cleared to '0', then a Controller Health Data Structure shall not be returned for Controllers associated with SR-IOV PFs.</p> <p>It is not an error if this bit is set to '1' and SR-IOV Physical <del>f</del>Functions do not exist.</p>
24	<p><b>Include PCI Functions (INCF):</b> <del>When-If</del> this bit is set to '1', then a Controller Health <del>Status-</del>Data Structure <del>is shall be</del> returned for Controllers associated with non-SR-IOV PCI Functions <del>unless excluded by other selection criteria as described in section 5.3.1.</del></p> <p>If this bit is cleared to '0', then a Controller Health Data Structure shall not be returned for Controllers associated with non-SR-IOV PCI Functions.</p> <p>It is not an error if this bit is set to '1' and non-SR-IOV PCI Functions do not exist.</p>

**Figure 77: Controller Health Status Poll – NVMe Management Dword 0**

Bits	Description
23:16	<p><b>Maximum Response Entries (MAXRENT):</b> This field specifies the maximum number of Controller Health Data Structure entries that may be returned in the completion. This is a 0's based field. The maximum number of entries is 255. If 256 entries are specified by this field, then an Invalid Parameter Error Response with the PEL field indicating this field shall be returned. <del>Specifying 256 entries is interpreted as an Invalid Parameter.</del></p>
15:00	<p><b>Starting Controller ID (SCTLID):</b> This field specifies the starting Controller ID <del>of the first Controller whose Controller Health Data Structure may be returned.</del></p> <p>If this field specifies a Controller ID that is less than or equal to the maximum Controller ID in the NVM Subsystem, then for each Controller in the NVM Subsystem:</p> <ul style="list-style-type: none"> <li>if the Controller ID of the Controller is greater than or equal to the value in this field, then the Controller's Controller Health Data Structure shall be returned unless excluded by other selection criteria as described in <a href="#">section 5.3.1</a>; or</li> <li>if the Controller ID of the Controller is less than the value in this field, then the Controller's Controller Health Data Structure shall not be returned.</li> </ul> <p>If this field specifies a Controller ID that is greater than the maximum Controller ID in the NVM Subsystem, then an Invalid Parameter Error Response with the PEL field indicating this field shall be returned.</p>

**Figure 78: Controller Health Status Poll – NVMe Management Dword 1**

Bits	Description
31	<p><b>Clear Changed Flags (CCF):</b> <del>When</del> If this bit is set to '1', then the Management Endpoint shall perform the following steps atomically in the order listed <del>the Controller Health Status Changed Flags are cleared in Controllers whose Controller Health Data Structure is contained in the Response Data:</del></p> <ol style="list-style-type: none"> <li>perform the selection criteria based on the Controller Health Status Changed Flags field as described in <a href="#">section 5.3.1.2</a>;</li> <li>copy the Controller Health Status Changed Flags field to the Controller Health Status Changed field in the Controller Health Data Structure; and</li> <li>clear each bit in the Controller Health Status Changed Flags field to '0' in Controllers whose Controller Health Data Structure is contained in the Response Data.</li> </ol> <p>If this bit is set to '1', then the following bits in the Controller Status field in the Controller Health Data Structure (refer to <a href="#">Figure 80</a>) shall be cleared to '0':</p> <ul style="list-style-type: none"> <li>Namespace Attribute Changed (NAC);</li> <li>Firmware Activated (FA); and</li> <li>Telemetry Controller-Initiated Data Available (TCIDA).</li> </ul> <p>The Controller Health Status Changed Flags field and the following bits in the Controller Status field in the Controller Health Data Structure shall not be modified in Controllers whose Controller Health Data Structure is not contained in the Response Data:</p> <ul style="list-style-type: none"> <li>Namespace Attribute Changed (NAC);</li> <li>Firmware Activated (FA); and</li> <li>Telemetry Controller-Initiated Data Available (TCIDA).</li> </ul> <p>If this bit is cleared to '0', then the Controller Health Status Changed Flags field and the following bits in the Controller Status field in the Controller Health Data Structure shall not be modified in any Controller:</p> <ul style="list-style-type: none"> <li>Namespace Attribute Changed (NAC);</li> <li>Firmware Activated (FA); and</li> <li>Telemetry Controller-Initiated Data Available (TCIDA).</li> </ul>
30:05	Reserved

Figure 78: Controller Health Status Poll – NVMe Management Dword 1

Bits	Description
04	<p><b>Critical Warning (CWARN):</b> <del>When</del>-If this bit is set to '1', then a Controller Health Data Structure <del>is</del>-shall be returned for Controllers with the Critical Warning bit set to '1' in their Controller Health Status Changed Flags field unless excluded by other selection criteria as described in <a href="#">section 5.3.1</a>.</p> <p>If this bit is set to '1', then a Controller Health Data Structure shall not be returned for Controllers with the Critical Warning bit cleared to '0' in their Controller Health Status Changed Flags field unless included by other selection criteria as described in <a href="#">section 5.3.1</a>.</p> <p>If this bit is cleared to '0', then the Critical Warning bit in the Controller Health Status Changed Flags field shall not be included in the selection criteria described in <a href="#">section 5.3.1</a>.</p>
03	<p><b>Available Spare (SPARE):</b> <del>When</del>-If this bit is set to '1', then a Controller Health Data Structure <del>is</del>-shall be returned for Controllers with the Available Spare bit set to '1' in their Controller Health Status Changed Flags field unless excluded by other selection criteria as described in <a href="#">section 5.3.1</a>.</p> <p>If this bit is set to '1', then a Controller Health Data Structure shall not be returned for Controllers with the Available Spare bit cleared to '0' in their Controller Health Status Changed Flags field unless included by other selection criteria as described in <a href="#">section 5.3.1</a>.</p> <p>If this bit is cleared to '0', then the Available Spare bit in the Controller Health Status Changed Flags field shall not be included in the selection criteria described in <a href="#">section 5.3.1</a>.</p>
02	<p><b>Percentage Used (PDLU):</b> <del>When</del>-If this bit is set to '1', then a Controller Health Data Structure <del>is</del>-shall be returned for Controllers with the Percent Used bit set to '1' in their Controller Health Status Changed Flags field unless excluded by other selection criteria as described in <a href="#">section 5.3.1</a>.</p> <p>If this bit is set to '1', then a Controller Health Data Structure shall not be returned for Controllers with the Percent Used bit cleared to '0' in their Controller Health Status Changed Flags field unless included by other selection criteria as described in <a href="#">section 5.3.1</a>.</p> <p>If this bit is cleared to '0', then the Percent Used bit in the Controller Health Status Changed Flags field shall not be included in the selection criteria described in <a href="#">section 5.3.1</a>.</p>
01	<p><b>Composite Temperature Changes (CTEMP):</b> <del>When</del>-If this bit is set to '1', then a Controller Health Data Structure <del>is</del>-shall be returned for Controllers with the Composite Temperature bit set to '1' in their Controller Health Status Changed Flags field unless excluded by other selection criteria as described in <a href="#">section 5.3.1</a>.</p> <p>If this bit is set to '1', then a Controller Health Data Structure shall not be returned for Controllers with the Composite Temperature bit cleared to '0' in their Controller Health Status Changed Flags field unless included by other selection criteria as described in <a href="#">section 5.3.1</a>.</p> <p>If this bit is cleared to '0', then the Composite Temperature bit in the Controller Health Status Changed Flags field shall not be included in the selection criteria described in <a href="#">section 5.3.1</a>.</p>
00	<p><b>Controller Status Changes (CSTS):</b> <del>When</del>-If this bit is set to '1', then a Controller Health Data Structure <del>is</del>-shall be returned for Controllers with the <del>Ready, Controller Fatal Status, Shutdown Status, NVM Subsystem Reset Occurred, Controller Enable Change Occurred, Namespace Attribute Changed, or Firmware Activated bit</del> Controller Status Change bit set to '1' in their Controller Health Status Changed Flags field unless excluded by other selection criteria as described in <a href="#">section 5.3.1</a>.</p> <p>If this bit is set to '1', then a Controller Health Data Structure shall not be returned for Controllers with the Controller Status Change bit cleared to '0' in their Controller Health Status Changed Flags field unless included by other selection criteria as described in <a href="#">section 5.3.1</a>.</p> <p>If this bit is cleared to '0', then the Controller Status Change bit in the Controller Health Status Changed Flags field shall not be included in the selection criteria described in <a href="#">section 5.3.1</a>.</p>

The Controller Health Status Poll Response Messages use the NVMe Management Response field with the format shown in Figure 79.

The Response Data field size may vary based on the number of Controllers selected using the selection criteria described in section 5.3.1 whose Controller Health Data Structure has changed and based on the number of Controllers whose Controller Health Data Structure is filtered out by Controller type (refer to section 5.3.1) or Controller Health Status Changed Flags (refer to section 5.3.2). The Response Entries field indicates the number of Controller Health Data Structures that are contained in the Response Data.

**Figure 79: Controller Health Status Poll – NVMe Management Response**

Bits	Description
23:16	<b>Response Entries (RENT):</b> This field specifies the number of Controller Health Data Structure Entries present in the Response Data for this Response Message. This is a 1-based field.
15:00	Reserved

The Controller Health Data Structure, shown in Figure 80, contains the health status attributes that are shall be tracked for each Controller. When the command is processed without error, health status is returned for up to 255 Controllers starting at or above the Starting Controller ID (SCTLID). Controllers are returned in ascending order of Controller Identifier. Up to 255 Controller Health Data Structures for Controllers with Controller IDs greater than or equal to the Starting Controller ID (SCTLID) shall be returned starting at or above offset 0h of the Response Data with no padding bytes between consecutive Controller Health Data Structures. Controller Health Data Structures shall be returned in ascending order based on their Controller IDs (i.e., the Controller Health Data Structure for the Controller with the lowest Controller ID that matches the section criteria in section 5.3.1 shall start at offset 0, the Controller Health Data Structure with the second lowest Controller ID that matches the section criteria in section 5.3.1 shall start at offset 16, etc.). The Response Data size shall be equal to the number of Controller Health Data Structures returned multiplied by the size of the Controller Health Data Structure.

**Figure 80: Controller Health Data Structure (CHDS)**

Bytes	Description
01:00	<b>Controller Identifier (CTLID):</b> This field shall indicate specifies the Controller Identifier of the Controller with which the data contained in this data structure is associated.

<b>Controller Status (CSTS):</b> This field reports the Controller status.		
Bits	Reset	Description
15:09 <del>08</del>	0	Reserved
08	Hwlnit	<p><b>Telemetry Controller-Initiated Data Available (TCIDA):</b> If the Telemetry Controller-Initiated Data Available field in the Telemetry Controller-Initiated log page (refer to the NVM Express Base Specification) transitions from 0h to 1h in this Controller, then this bit shall be set to '1'.</p> <p>If this bit is set to '1', then the Telemetry Controller-Initiated log page shall contain saved internal controller state or saved internal NVM Subsystem state available by issuing a Get Log Page command using the out-of-band mechanism. If this bit is cleared to '0' and the Telemetry Controller-Initiated log page is Controller in scope, then the Telemetry Controller-Initiated log page shall not contain saved internal controller state but may contain saved NVM Subsystem state (refer to the CCSF.TCIDA bit) available by issuing a Get Log Page command using the out-of-band mechanism.</p> <p>If a Controller Health Status Poll command is processed with the Clear Changed Flags bit set to '1', then this bit shall be cleared to '0'.</p> <p>The value of this field shall persist across all resets and power cycles.</p>
07	Hwlnit	<p><b>Firmware Activated (FA):</b> <del>This bit is set to '1' when</del> If a new firmware image is activated, then this bit shall be set to '1'. Firmware activation is described in the NVM Express Base Specification.</p> <p><del>The reset value of this bit is set to '1' if</del> a reset caused a new firmware image to be activated, then the reset value of this bit shall be '1'.</p> <p>If a Controller Health Status Poll command is processed with the Clear Changed Flags bit set to '1', then this bit shall be cleared to '0'.</p>
06	0	<p><b>Namespace Attribute Changed (NAC):</b> This bit <del>shall be is</del> set to '1' under the same conditions that causes the Namespace Attribute Changed asynchronous event to be sent <del>with the exception that if</del> Namespace Attribute Notices are <del>not required to be</del> enabled <del>as specified in</del> (refer to the NVM Express Base Specification). <del>This bit may be set to '1' regardless of whether Namespace Attribute Notices are enabled or not.</del></p> <p>If a Controller Health Status Poll command is processed with the Clear Changed Flags bit set to '1', then this bit shall be cleared to '0'.</p>
05	0	<p><b>Controller Enable Change Occurred (CECO):</b> This bit is set to '1' when the Enable bit (refer to CC.EN in the NVM Express Base Specification) changes state.</p>
04	Hwlnit	<p><b>NVM Subsystem Reset Occurred (NSSRO):</b> <del>This bit corresponds to the value of the NVM Subsystem Reset Occurred (refer to CSTS.NSSRO in the NVM Express Base Specification) bit.</del></p> <p>If an NVM Subsystem Reset occurs due to any reason other than application of main power and does not cause activation of a new firmware image, then this bit shall be set to '1'.</p> <p>If an NVM Subsystem Reset occurs due to application of main power or causes activation of a new firmware image, then this bit shall be cleared to '0'.</p>
03:02	00b	<p><b>Shutdown Status (SHST):</b> This field <del>shall indicate corresponds to</del> the value of the Shutdown Status field (refer to CSTS.SHST in the NVM Express Base Specification) <del>field.</del></p>
01	Hwlnit	<p><b>Controller Fatal Status (CFS):</b> This bit <del>shall indicate corresponds to</del> the value of the Controller Fatal Status bit (refer to CSTS.CFS in the NVM Express Base Specification) <del>bit.</del></p>

**Figure 80: Controller Health Data Structure (CHDS)**

Bytes	Description		
	00	0	<b>Ready (RDY):</b> This bit shall indicate <del>corresponds to</del> the value of the Ready bit (refer to CSTS.RDY in the NVM Express Base Specification) <del>bit</del> .
05:04			<b>Composite Temperature (CTEMP):</b> This field indicates <del>contains</del> a value corresponding to a temperature in Kelvins that represents the current composite temperature of the Controller and Namespace(s) associated with that Controller. The value of this field shall indicate <del>corresponds to</del> the value of the <b>Composite Temperature</b> field in the Controller's SMART / Health Information <a href="#">Log page</a> .
06			<b>Percentage Used (PDLU):</b> This field indicates <del>an implementation-specific contains a vendor specific</del> estimate of the percentage of NVM Subsystem life used based on the actual usage and the manufacturer's prediction of NVM <b>Subsystem</b> life. The value of this field shall indicate <del>corresponds to</del> the value of the <b>Percent Used</b> field in the Controller's SMART / Health Information <a href="#">Log page</a> .
07			<b>Available Spare (SPARE):</b> This field indicates <del>contains</del> a normalized percentage (0% to 100%) of the remaining spare capacity available. The value of this field shall indicate <del>corresponds to</del> the value of the <b>Available Spare</b> field in the Controller's SMART / Health Information <a href="#">Log page</a> .
08			<b>Critical Warning (CWARN):</b> This field indicates critical warnings for <del>the state of</del> the Controller. The value of this field shall indicate <del>corresponds to</del> the value of the <b>Critical Warning</b> field in the Controller's SMART / Health Information <a href="#">Log page</a> .
	Bits	Description	
	7:6	Reserved	
	5	<b>Persistent Memory Region Error (PMRE):</b> This bit shall indicate the same value as bit 5 in the <b>Critical Warning</b> field in the Controller's SMART / Health Information <a href="#">log page</a> <del>is set to '1' when the Persistent Memory Region has become read-only or unreliable.</del>	
	4	<b>Volatile Memory Backup Failed (VMBF):</b> This bit shall indicate the same value as bit 4 in the <b>Critical Warning</b> field in the Controller's SMART / Health Information <a href="#">log page</a> <del>is set to '1' when the volatile memory backup device has failed.</del>	
	3	<b>Read Only (RO):</b> This bit shall indicate the same value as bit 3 in the <b>Critical Warning</b> field in the Controller's SMART / Health Information <a href="#">log page</a> <del>is set to '1' when the media has been placed in read-only mode.</del>	
	2	<b>Reliability Degraded (RD):</b> This bit shall indicate the same value as bit 2 in the <b>Critical Warning</b> field in the Controller's SMART / Health Information <a href="#">log page</a> <del>is set to '1' when NVM Subsystem reliability has been degraded due to significant media related errors or an internal error.</del>	
	1	<b>Temperature Above or Under Threshold (TAUT):</b> This bit shall indicate the same value as bit 1 in the <b>Critical Warning</b> field in the Controller's SMART / Health Information <a href="#">log page</a> <del>is set to '1' when a temperature is above an over temperature threshold or below an under temperature threshold.</del>	
	0	<b>Spare Threshold (ST):</b> This bit shall indicate the same value as bit 0 in the <b>Critical Warning</b> field in the Controller's SMART / Health Information <a href="#">log page</a> <del>is set to '1' when the available spare has fallen below the available spare threshold.</del>	
10:09			<b>Controller Health Status Changed (CHSC):</b> This field shall indicate the value of the Controller Health Status Changed Flags field (refer to <a href="#">Figure 81</a> ).
15:11 <del>09</del>			Reserved

Associated with each Controller in the NVM Subsystem is a set of Controller Health Status Changed Flags shown in Figure 81. The Controller Health Status Changed Flags ~~are shall be~~ set as described in [Figure 81](#) ~~if when~~ the corresponding field/bit in the Controller Health Data Structure changes ~~state as described in Figure 84~~. Figure 82 shows a graphical representation of which field(s)/bit(s) in the Controller Health Data Structure ~~are shall be~~ associated with each bit in the Controller Health Status Changed Flags. ~~When~~ If 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 Flags ~~is shall~~ also ~~be~~ set to '1'. The Controller Health Status Changed Flags ~~are shall be~~ cleared to 0h after the Controller selection criteria has been evaluated as described in [section 5.3.1](#) in any Controllers whose Controller Health Data Structure is returned in the Success Response to a Controller Health Status Poll Command Message with the Clear Changed Flags bit set to '1'.

### 5.3.1 Controller Selection Criteria

A Controller Health Status Poll response ~~may~~ returns the Controller Health Data Structure for up to 255 Controllers in the Response Data field. An NVM Subsystem ~~may~~ contains up to 64 Ki Controllers, so a method is required to limit the size of the Response Message. The Starting Controller ID field in the Command Message specifies the Controller ID of the first Controller whose Controller Health Data Structure may be returned in the Response Data field. The Maximum Response Entries field specifies the maximum number of Controllers whose Controller Health Data Structure 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, where M is equal to the Maximum Response Entries field and N is equal to the Starting Controller ID field.~~ The Response Data field shall contain the ~~entire~~ Controller Health Status Data Structure for the first M Controllers in order of ascending Controller ID, where M is equal to the value in the Maximum Response Entries field, for ~~all~~ any Controllers that:

- has a Controller ID that is greater than or equal to the value in the Starting Controller ID field;
- ~~do not match~~ matches the ~~filtering controller type selection~~ criteria in ~~Controller Health Status Poll – NVMe Management Dword 0~~ (refer to [section 5.3.1.1](#)); and
- either:
  - a. ~~and that have one or more Controller Health Status Changed Flags that are: a) set and b) do not match~~ matches the ~~filtering~~ Controller Health Status Changed Flags field selection criteria in ~~Controller Health Status Poll – NVMe Management Dword 1~~ (refer to [section 5.3.1.2](#)); or
  - b. has been requested to report all changes to the Controller's Controller Health Status Changed Flags field (i.e., the Report All bit is set to '1' in the NVMe Management Dword 0 field in the Controller Health Status Poll command).

The Response Data field shall not contain the Controller Health ~~Status~~ Data Structure for any Controllers that ~~do not~~ meet the ~~filtering selection~~ criteria in ~~this section sections 5.3.1 or 5.3.2~~.

#### 5.3.1.1 ~~Filtering Selection Criteria~~ by Controller Type

The ~~Controllers whose~~ Controller Health Data Structures ~~that~~ are returned by the Controller Health Status Poll ~~may be filtered~~ are selected (i.e., ~~excluded from being included in the Response Data field regardless of the state of the Controller Health Status Changed Flags~~) by based on Controller type (i.e., non-SR-IOV PCI Function, SR-IOV PF, and SR-IOV VF). Controller type ~~filtering selection~~ is controlled by the Include PCI Functions (INCF), Include SR-IOV PFs (INCPF), and Include SR-IOV VFs (INCVF) bits fields in the NVMe Management Dword 0 field. ~~When~~ If one or more of these ~~bits~~ (INCF, INCPF, or INCVF) bits ~~are is~~ set to '1', then a Controller Health Data Structures for Controllers corresponding to that type of PCI Function ~~are shall be~~ included in the Response Data field ~~unless excluded by other selection criteria as described in section 5.3.1; else, the Controller Health Data Structure for that Controller is excluded from the Response Data field.~~

#### 5.3.1.2 ~~Filtering Selection Criteria~~ by Controller Health Status Changed Flags

If the Report All bit is cleared to '0', then ~~the~~ the Controllers whose Controller Health Data Structures ~~that~~ are returned by the Controller Health Status Poll ~~may are~~ also ~~be filtered by~~ selected based on the Controller

Health Status Changed Flags. ~~Filtering Selection of Controllers changes~~ by Controller Health Status Changed Flags is controlled by ~~some of the~~ CWARN, SPARE, PDLU, CTEMP, and CSTS bits in the NVMe Management Dword 1 field (refer to **Figure 78**). ~~When~~ If one or more of ~~these the~~ CWARN, SPARE, PDLU, CTEMP, or CSTS bits in the NVMe Management Dword 1 field are set to '1' and any of the corresponding bit(s) in the Controller Health Status Changed Flags for the Controller are also set to '1' (refer to Figure 78 for Controller Health Status Changed Flags associated with each bit in the NVMe Management Dword 1 field), then the entire Controller Health Data Structure for that Controller ~~is shall be~~ returned in the Response Data field ~~unless excluded by other selection criteria as described in section 5.3.1; else, the Controller Health Data Structure for that Controller is excluded from the Response Data field.~~

If the Report All bit is cleared to '0', then ~~the~~ the contents returned in the Controller Health Data Structure for ~~filtered~~ the CWARN, SPARE, PDLU, CTEMP, and CSTS fields are undefined if the corresponding CWARN, SPARE, PDLU, CTEMP, or CSTS bit in the NVMe Management Dword 1 field is cleared to '0'. If the Report All bit is set to '1', then the contents returned in the Controller Health Data Structure for the CWARN, SPARE, PDLU, CTEMP, and CSTS fields shall be valid regardless of the value of the corresponding CWARN, SPARE, PDLU, CTEMP, or CSTS bit in the NVMe Management Dword 1 field.

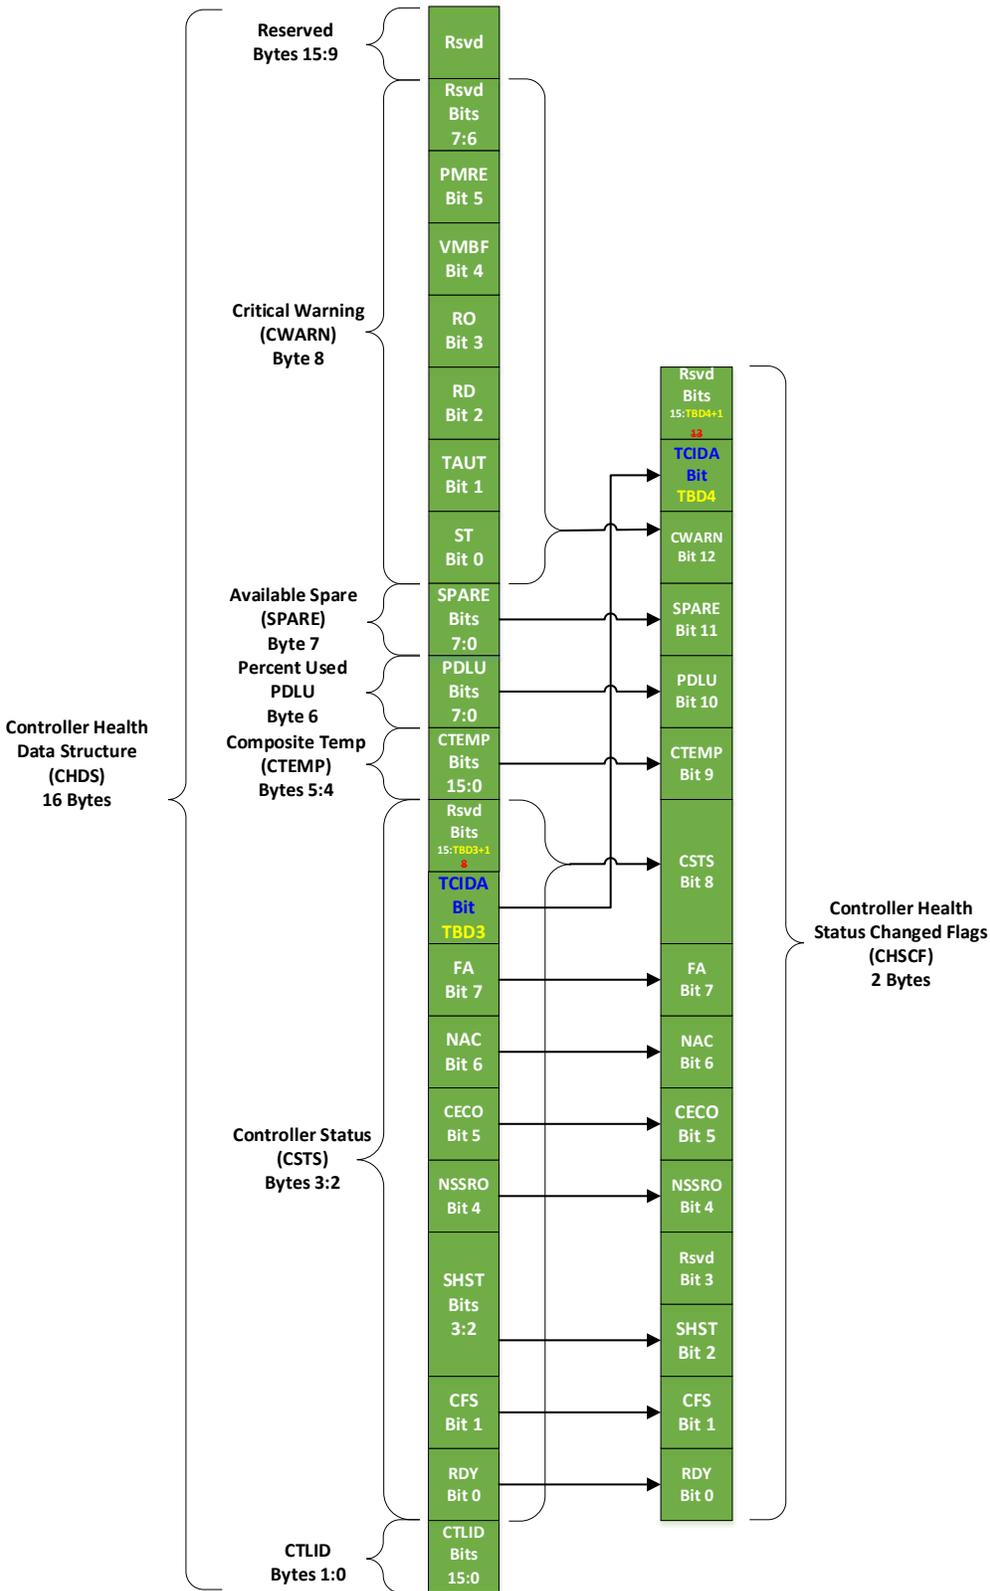
**Figure 81: Controller Health Status Changed Flags (CHSCF)**

Bits	Reset	Description
15:14 <del>13</del>	0	Reserved
13	HwInit	<b>Telemetry Controller-Initiated Data Available (TCIDA):</b> If the Telemetry Controller-Initiated Data Available bit in the Controller Health Data Structure transitions from '0' to '1', then this bit shall be set to '1'. The value of this field shall persist across all resets and power cycles.
12	0	<b>Critical Warning (CWARN):</b> <del>This bit is set to '1' when</del> If any of the Critical Warning field bits in the Controller Health Data Structure <del>changes state, then this bit shall be set to '1'. transitions from '0' to '1'.</del>
11	0	<b>Available Spare (SPARE):</b> <del>This bit is set to '1' when</del> If the Available Spare field in the Controller Health Data Structure changes state, <del>then this bit shall be set to '1'.</del>
10	0	<b>Percentage Used (PDLU):</b> <del>This bit is set to '1' when</del> If the Percentage Used field in the Controller Health Data Structure changes state, <del>then this bit shall be set to '1'.</del>
09	0	<b>Composite Temperature Change (CTEMP):</b> <del>This bit is set to '1' when</del> If the Composite Temperature field in the Controller Health Data Structure changes state, <del>then this bit shall be set to '1'.</del>
08	HwInit	<b>Controller Status Change (CSTS):</b> <del>This bit is set to '1' when</del> If any bit or field in the Controller Status field in the Controller Health Data Structure (e.g., the Shutdown Status field, <del>in the Controller Health Data Structure changes state or when</del> the Ready bit, the Controller Fatal Status bit, the NVM Subsystem Reset Occurred bit, the Controller Enable Change Occurred bit, the Namespace Attribute Changed bit, <del>or the</del> Firmware Activated bit, or the Telemetry Controller-Initiated Data Available bit) changes state, then this bit shall be set to '1' <del>in the Controller Health Data Structure transitions from '0' to '1'.</del>
07	HwInit	<b>Firmware Activated (FA):</b> <del>This bit is set to '1' when</del> If the Firmware Activated bit in the Controller Health Data Structure transitions from '0' to '1', <del>then this bit shall be set to '1'.</del>
06	0	<b>Namespace Attribute Changed (NAC):</b> <del>This bit is set to '1' when</del> If the Namespace Attribute Changed bit in the Controller Health Data Structure transitions from '0' to '1', <del>then this bit shall be set to '1'.</del>
05	0	<b>Controller Enable Change Occurred (CECO):</b> <del>This bit is set to '1' when</del> If the Controller Enable Change Occurred bit in the Controller Health Data Structure changes state, <del>then this bit shall be set to '1' transitions from '0' to '1'.</del>

**Figure 81: Controller Health Status Changed Flags (CHSCF)**

Bits	Reset	Description
04	HwInit	<b>NVM Subsystem Reset Occurred (NSSRO):</b> <del>This bit is set to '1' when</del> If the NVM Subsystem Reset Occurred bit in the Controller Health Data Structure transitions from '0' to '1', <del>then this bit shall be set to '1'.</del>
03	0	Reserved
02	0	<b>Shutdown Status (SHST):</b> <del>This bit is set to '1' when</del> If the Shutdown Status field in the Controller Health Data Structure changes state, <del>then this bit shall be set to '1'.</del>
01	HwInit	<b>Controller Fatal Status (CFS):</b> <del>This bit is set to '1' when</del> If the Controller Fatal Status bit in the Controller Health Data Structure <del>changes state, then this bit shall be set to '1'</del> <del>transitions from '0' to '1'.</del>
00	0	<b>Ready (RDY):</b> <del>This bit is set to '1' when</del> If the Ready bit in the Controller Health Data Structure <del>changes state, then this bit shall be set to '1'</del> <del>transitions from '0' to '1'.</del>

**Figure 82: Controller Health Data Structure to Controller Health Status Changed Flags Mapping**



**Modify section 5.6 as follows:**

## 5.6 NVM Subsystem Health Status Poll

The NVM Subsystem Health Status Poll command is used to efficiently determine changes in health status attributes associated with the NVM Subsystem.

The NVM Subsystem Health Status Poll command operates independently using the out-of-band mechanism and the in-band tunneling mechanism.

An NVMe Storage Device or NVMe Enclosure supporting the NVM Subsystem Health Status Poll command using the out-of-band mechanism shall have an independent copy of the NVM Subsystem Health Data Structure (refer to Figure 90) dedicated to the out-of-band mechanism. In the out-of-band mechanism, an NVM Subsystem Health Status Poll command only applies to the copy of the NVM Subsystem Health Data Structure dedicated to the out-of-band mechanism.

An NVMe Storage Device or NVMe Enclosure supporting the NVM Subsystem Health Status Poll command using the in-band tunneling mechanism shall have an independent copy of the NVM Subsystem Health Data Structure dedicated to the in-band tunneling mechanism. In the in-band tunneling mechanism, an NVM Subsystem Health Status Poll command only applies to the copy of the NVM Subsystem Health Data Structure dedicated to the in-band tunneling mechanism.

The NVM Subsystem Health Status Poll command uses NVMe Management Dword 1 as shown in Figure 89.

**Figure 89: NVM Subsystem Health Status Poll - NVMe Management Dword 1**

Bits	Description
31	<p><b>Clear Status (CS):</b> <del>When</del> If this bit is set to '1', then the Management Endpoint shall perform the following steps atomically in the order listed:</p> <ol style="list-style-type: none"><li>copy the current value of the <del>state of reported</del> Composite Controller Status Flags field (refer to Figure TBD2) to the Composite Controller Status field of the Response Message (refer to Figure 90); and</li><li><del>is cleared</del> the Composite Controller Status Flags field to 0h.</li></ol> <p>If this bit is cleared to '0', then the Management Endpoint shall copy the current value of the Composite Controller Status Flags field (refer to Figure TBD2) to the Composite Controller Status field of the Response Message (refer to Figure 90) and shall not modify the Composite Controller Status Flags field.</p>
30:00	Reserved

All other ~~command-specific~~ ~~command-specific~~ fields are reserved.

The NVM Subsystem Health Data Structure, shown in Figure 90, ~~is shall be~~ returned starting at offset 0h in the Response Data of a Successful Response ~~Message~~. The NVM Subsystem Health Status Poll ~~Command~~ ~~command~~ ~~Response Messages~~ ~~responses~~ do not use the NVMe Management Response field and this field ~~is shall be~~ reserved. The Response Data field ~~contains the NVM Subsystem Health Data Structure and is always shall be~~ the size of the NVM Subsystem Health Data Structure.

Figure TBD2: Composite Controller Status Data Structure (CCSDS)

Bytes	Description
-------	-------------

1:0	<p><b>Composite Controller Status Flags (CCSF):</b> This field indicates the composite status of all Controllers in the NVM Subsystem.</p> <p>The bits in this field are cleared to '0' as described in the Clear Status field (refer to Figure 89) after the NVM Subsystem Health Data Structure (refer to Figure 90) is returned in a Success Response associated with an NVM Subsystem Health Status Poll command where the Clear Status bit set. A Configuration Set command that selects specifies a Configuration Identifier value of 02h (Health Status Change) in the NVMe Management Dword 1 field may be used to clear selected bits to '0' (refer to section 5.1.2).</p>		
	Bits	Reset <sup>1</sup>	Description
	15:143	0	Reserved
	13	Hwlnit	<p><b>Telemetry Controller-Initiated Data Available (TCIDA):</b> If the Telemetry Controller-Initiated Data Available bit in the Controller Health Status Changed Flags field transitions from '0' to '1' in one or more Controllers in the NVM Subsystem, then this bit shall be set to '1'.</p> <p>If the Telemetry Controller-Initiated log page is NVM Subsystem in scope, then:</p> <ul style="list-style-type: none"> <li>if this bit is set to '1', then the Telemetry Controller-Initiated log page shall contain saved internal NVM Subsystem state available by issuing a Get Log Page command to any Controller in the NVM Subsystem using the out-of-band mechanism; and</li> <li>if this bit is cleared to '0', then the Telemetry Controller-Initiated log page shall not contain saved internal NVM Subsystem state available by issuing a Get Log Page command to any Controller in the NVM Subsystem using the out-of-band mechanism.</li> </ul> <p>If the Telemetry Controller-Initiated log page is Controller in scope, then refer to the Telemetry Controller-Initiated Data Available bit in the Controller Status field of each Controller in the NVM Subsystem to determine the availability of the Telemetry Controller-Initiated log page.</p> <p>The value of this field shall persist across all resets and power cycles.</p>
	12	0	<p><b>Critical Warning (CWARN):</b> <del>This bit is set to '1' when</del> If the Critical Warning bit in the Controller Health Status Changed Flags field transitions from '0' to '1' in one or more Controllers in the NVM Subsystem, then this bit shall be set to '1'.</p>
	11	0	<p><b>Available Spare (SPARE):</b> <del>This bit is set to '1' when</del> If the Available Spare bit in the Controller Health Status Changed Flags field transitions from '0' to '1' in one or more Controllers in the NVM Subsystem, then this bit shall be set to '1'.</p>
	10	0	<p><b>Percentage Used (PDLU):</b> <del>This bit is set to '1' when</del> If the Percentage Used field bit in the Controller Health Status Changed Flags field transitions from '0' to '1' in one or more Controllers in the NVM Subsystem, then this bit shall be set to '1'.</p>
	09	0	<p><b>Composite Temperature Change (CTEMP):</b> <del>This bit is set to '1' when</del> If the Composite Temperature field bit in the Controller Health Status Changed Flags field transitions from '0' to '1' in one or more Controllers in the NVM Subsystem, then this bit shall be set to '1'.</p>
	08	Hwlnit	<p><b>Controller Status Change (CSTS):</b> <del>This bit is set to '1' when</del> If the Controller Status Change field bit in the Controller Health Status Changed Flags field transitions from '0' to '1' in one or more Controllers in the NVM Subsystem, then this bit shall be set to '1'.</p>
	07	Hwlnit	<p><b>Firmware Activated (FA):</b> <del>This bit is set to '1' when</del> If the Firmware Activated bit in the Controller Health Status Changed Flags field transitions from '0' to '1' in one or more Controllers in the NVM Subsystem, then this bit shall be set to '1'.</p>
06	0	<p><b>Namespace Attribute Changed (NAC):</b> <del>This bit is set to '1' when</del> If the Namespace Attribute Changed bit in the Controller Health Status Changed Flags field transitions from '0' to '1' in one or more Controllers in the NVM Subsystem, then this bit shall be set to '1'.</p>	

	05	0	<b>Controller Enable Change Occurred (CECO):</b> <del>This bit is set to '1' when</del> If the Controller Enable Change Occurred bit in the Controller Health Status Changed Flags field transitions from '0' to '1' in one or more Controllers in the NVM Subsystem, then this bit shall be set to '1'.
	04	HwInit	<b>NVM Subsystem Reset Occurred (NSSRO):</b> <del>This bit is set to '1' when</del> If the value of the NVM Subsystem Reset Occurred ( <del>GSTS.NSSRO</del> ) bit in the Controller Health Status Changed Flags field transitions from a '0' to a '1' in one or more Controllers in the NVM Subsystem, then this bit shall be set to '1'.
	03	0	Reserved
	02	0	<b>Shutdown Status (SHST):</b> <del>This bit is set to '1' when</del> If the Shutdown Status bit in the Controller Health Status Changed Flags field transitions from '0' to '1' in one or more Controllers in the NVM Subsystem, then this bit shall be set to '1'.
	01	HwInit	<b>Controller Fatal Status (CFS):</b> <del>This bit is set to '1' when</del> If the Controller Fatal Status bit in the Controller Health Status Changed Flags field transitions from '0' to '1' in one or more Controllers in the NVM Subsystem, then this bit shall be set to '1'.
	00	0	<b>Ready (RDY):</b> <del>This bit is set to '1' when</del> If the Ready bit in the Controller Health Status Changed Flags field transitions from '0' to '1' in one or more Controllers in the NVM Subsystem, then this bit shall be set to '1'.

**NOTES:**

1. An NVM Subsystem Reset shall reset the instance of the Composite Controller Status Flags field ~~Controller Status bits~~ dedicated to the out-of-band mechanism and the instance of the Composite Controller Status Flags field ~~Controller Status bits~~ dedicated to each Controller in the NVM Subsystem.  
The instance of the Composite Controller Status Flags field ~~Controller Status bits~~ dedicated to a Controller shall be reset by a Controller Level Reset (refer to the NVM Express Base Specification) of that Controller. Note that a Controller Level Reset may affect the Composite Controller Status Flags field ~~Controller Status bits~~ in the out-of-band mechanism (e.g., a Controller Level Reset causes the CECO bit in the instance of the Composite Controller Status Flags field ~~Controller Status bits~~ dedicated to the out-of-band mechanism to be set to '1').  
No instance of the Composite Controller Status Flags field ~~Controller Status bits~~ shall be reset by any other resets other than the resets documented by this note.

Figure 90: NVM Subsystem Health Data Structure (NSHDS)

Bytes	Description																
0	<p><b>NVM Subsystem Status (NSS):</b> This field indicates the status of the NVM Subsystem.</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>7:6</td> <td>Reserved</td> </tr> <tr> <td>6</td> <td><b>Sanitize Failure Mode (SFM):</b> If the NVM Subsystem is in the Sanitize failure mode, then this bit shall be set to '1'. If the NVM Subsystem is not in the Sanitize failure mode, then this bit shall be cleared to '0'. The NVM Subsystem is in the Sanitize failure mode when the most recent sanitize operation failed and no recovery action has been completed successfully (refer to the NVM Express Base Specification).</td> </tr> <tr> <td>5</td> <td><b>Drive Functional (DF):</b> If the NVM Subsystem is functional, then this bit shall be set to '1' to indicate an NVM Subsystem is functional. If cleared to '0', then there is an unrecoverable failure detected in the NVM Subsystem, then this bit shall be cleared to '0'.</td> </tr> <tr> <td>4</td> <td><b>Reset Not Required (RNR):</b> This bit is set to '1' to indicate the NVM Subsystem does not require 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><b>Port 0 PCIe Link Active (P0LA):</b> This bit is set to '1' to indicate the first port's PCIe link on the port with the lowest Port Identifier is active up (i.e., the Data Link Control and Management State Machine is in the DL_Active state as defined by the PCI Express Base Specification), then this bit shall be set to '1'.  If the PCIe link on the port with the lowest Port Identifier is not active (i.e., the Data Link Control and Management State Machine is not in the DL_Active state as defined by the PCI Express Base Specification), then this bit shall be cleared to '0'. If cleared to '0', then the PCIe link is down.</td> </tr> <tr> <td>2</td> <td><b>Port 1 PCIe Link Active (P1LA):</b> This bit is set to '1' to indicate the first port's PCIe link on the port with the second lowest Port Identifier is active up (i.e., the Data Link Control and Management State Machine is in the DL_Active state as defined by the PCI Express Base Specification), then this bit shall be set to '1'.  If the PCIe link on the port with the second lowest Port Identifier is not active (i.e., the Data Link Control and Management State Machine is not in the DL_Active state as defined by the PCI Express Base Specification) or there is no port 1, then this bit shall be cleared to '0'. 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> </tbody> </table>	Bits	Description	7:6	Reserved	6	<b>Sanitize Failure Mode (SFM):</b> If the NVM Subsystem is in the Sanitize failure mode, then this bit shall be set to '1'. If the NVM Subsystem is not in the Sanitize failure mode, then this bit shall be cleared to '0'. The NVM Subsystem is in the Sanitize failure mode when the most recent sanitize operation failed and no recovery action has been completed successfully (refer to the NVM Express Base Specification).	5	<b>Drive Functional (DF):</b> If the NVM Subsystem is functional, then this bit shall be set to '1' to indicate an NVM Subsystem is functional. If cleared to '0', then there is an unrecoverable failure detected in the NVM Subsystem, then this bit shall be cleared to '0'.	4	<b>Reset Not Required (RNR):</b> This bit is set to '1' to indicate the NVM Subsystem does not require 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	<b>Port 0 PCIe Link Active (P0LA):</b> This bit is set to '1' to indicate the first port's PCIe link on the port with the lowest Port Identifier is active up (i.e., the Data Link Control and Management State Machine is in the DL_Active state as defined by the PCI Express Base Specification), then this bit shall be set to '1'.  If the PCIe link on the port with the lowest Port Identifier is not active (i.e., the Data Link Control and Management State Machine is not in the DL_Active state as defined by the PCI Express Base Specification), then this bit shall be cleared to '0'. If cleared to '0', then the PCIe link is down.	2	<b>Port 1 PCIe Link Active (P1LA):</b> This bit is set to '1' to indicate the first port's PCIe link on the port with the second lowest Port Identifier is active up (i.e., the Data Link Control and Management State Machine is in the DL_Active state as defined by the PCI Express Base Specification), then this bit shall be set to '1'.  If the PCIe link on the port with the second lowest Port Identifier is not active (i.e., the Data Link Control and Management State Machine is not in the DL_Active state as defined by the PCI Express Base Specification) or there is no port 1, then this bit shall be cleared to '0'. If cleared to '0', then the second port's PCIe link is down or not present.	1:0	Reserved
	Bits	Description															
	7:6	Reserved															
	6	<b>Sanitize Failure Mode (SFM):</b> If the NVM Subsystem is in the Sanitize failure mode, then this bit shall be set to '1'. If the NVM Subsystem is not in the Sanitize failure mode, then this bit shall be cleared to '0'. The NVM Subsystem is in the Sanitize failure mode when the most recent sanitize operation failed and no recovery action has been completed successfully (refer to the NVM Express Base Specification).															
	5	<b>Drive Functional (DF):</b> If the NVM Subsystem is functional, then this bit shall be set to '1' to indicate an NVM Subsystem is functional. If cleared to '0', then there is an unrecoverable failure detected in the NVM Subsystem, then this bit shall be cleared to '0'.															
	4	<b>Reset Not Required (RNR):</b> This bit is set to '1' to indicate the NVM Subsystem does not require 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	<b>Port 0 PCIe Link Active (P0LA):</b> This bit is set to '1' to indicate the first port's PCIe link on the port with the lowest Port Identifier is active up (i.e., the Data Link Control and Management State Machine is in the DL_Active state as defined by the PCI Express Base Specification), then this bit shall be set to '1'.  If the PCIe link on the port with the lowest Port Identifier is not active (i.e., the Data Link Control and Management State Machine is not in the DL_Active state as defined by the PCI Express Base Specification), then this bit shall be cleared to '0'. If cleared to '0', then the PCIe link is down.															
2	<b>Port 1 PCIe Link Active (P1LA):</b> This bit is set to '1' to indicate the first port's PCIe link on the port with the second lowest Port Identifier is active up (i.e., the Data Link Control and Management State Machine is in the DL_Active state as defined by the PCI Express Base Specification), then this bit shall be set to '1'.  If the PCIe link on the port with the second lowest Port Identifier is not active (i.e., the Data Link Control and Management State Machine is not in the DL_Active state as defined by the PCI Express Base Specification) or there is no port 1, then this bit shall be cleared to '0'. If cleared to '0', then the second port's PCIe link is down or not present.																
1:0	Reserved																
1	<p><b>Smart Warnings (SW):</b> This field indicates <del>contains</del> the inverted value of the Critical Warning field (byte 0) of the NVMe SMART / Health Information log page. Each bit in this field <del>is</del> shall be inverted from the value in the Critical Warning field of the SMART / Health Information log page as defined by the NVM Express Base Specification <del>definition (i.e., the management interface shall indicate a '0' value while the corresponding bit is set to '1' in the SMART / Health Information log page. Refer to the NVM Express Base Specification for bit definitions.</del></p> <p>If there are multiple Controllers in the NVM Subsystem, the Responder 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"> <li><del>C</del>leared to '0' if the corresponding bit in the Critical Warning field of the SMART / Health Information log page of any Controller in the NVM Subsystem is set to '1' <del>any Controller in the NVM Subsystem indicates a critical warning for that corresponding bit;</del> or</li> <li><del>S</del>set to '1' if the corresponding bit in the Critical Warning field of the SMART / Health Information log page in all Controllers in the NVM Subsystem is cleared to '0' <del>all Controllers in the NVM Subsystem do not indicate a critical warning for the corresponding bit.</del></li> </ul>																

**Figure 90: NVM Subsystem Health Data Structure (NSHDS)**

Bytes	Description																
2	<p><b>Composite Temperature (CTEMP):</b> This field indicates <del>the current</del> information related to the composite temperature of the NVM Subsystem <del>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.</del> The composite temperature of the NVM Subsystem shall be calculated at least every 5 s as follows:</p> <ul style="list-style-type: none"> <li>if there are one or more Controllers in the NVM Subsystem with a Composite Temperature in the SMART / Health Information log page that is less than or equal to an under-temperature threshold; and</li> <li>no Controllers in the NVM Subsystem with a Composite Temperature in the SMART / Health Information log page that is greater than or equal to an over-temperature threshold (refer to the NVM Express Base Specification),</li> </ul> <p>then the composite temperature of the NVM Subsystem shall be the same temperature as the Composite Temperature from the SMART / Health Information log page of the coldest Controller in the NVM Subsystem; otherwise, the composite temperature of the NVM Subsystem shall be the same temperature as the Composite Temperature from the SMART / Health Information log page of the hottest Controller in the NVM Subsystem.</p> <p>The reported temperature range is <del>vendor implementation</del> specific <del>and shall not exceed the range -60 °C to +127 °C.</del> The <del>8-bit format of values for this field are as follows: the data is shown below.</del></p> <p><del>This field should not report a temperature that is older than 1 s. If recent data is not available, the Responder should indicate a value of 80h for this field.</del></p> <table border="1" data-bbox="344 907 1360 1495"> <thead> <tr> <th data-bbox="344 907 516 940">Value</th> <th data-bbox="516 907 1360 940">Description</th> </tr> </thead> <tbody> <tr> <td data-bbox="344 940 516 1054">00h to 7Eh</td> <td data-bbox="516 940 1360 1054"><del>Temperature is measured in degrees Celsius (If the composite temperature of the NVM Subsystem is greater than or equal to 0 °C and less than or equal to 126 °C),</del> then this field shall indicate the composite temperature of the NVM Subsystem in degrees Celsius.</td> </tr> <tr> <td data-bbox="344 1054 516 1108">7Fh</td> <td data-bbox="516 1054 1360 1108">If the composite temperature of the NVM Subsystem is greater than or equal to 127 °C <del>or higher,</del> then this field shall indicate a value of 7Fh.</td> </tr> <tr> <td data-bbox="344 1108 516 1192">80h</td> <td data-bbox="516 1108 1360 1192"><del>No temperature data or temperature data</del> If the composite temperature of the NVM Subsystem is <del>more the</del> greater than 5 s old, then this field shall indicate a value of 80h.</td> </tr> <tr> <td data-bbox="344 1192 516 1276">81h</td> <td data-bbox="516 1192 1360 1276">If the composite temperature of the NVM Subsystem is not accurate due to the failure of one or more temperature sensors, then this field shall indicate a value of 81h. <del>Temperature sensor failure</del></td> </tr> <tr> <td data-bbox="344 1276 516 1306">82h to C3h</td> <td data-bbox="516 1276 1360 1306">Reserved</td> </tr> <tr> <td data-bbox="344 1306 516 1360">C4h</td> <td data-bbox="516 1306 1360 1360">If the composite temperature of the NVM Subsystem is less than or equal to <del>Temperature is</del> -60 °C <del>or lower,</del> then this field shall indicate a value of C4h.</td> </tr> <tr> <td data-bbox="344 1360 516 1495">C5h to FFh</td> <td data-bbox="516 1360 1360 1495">If the composite temperature of the NVM Subsystem is less than or equal to -1 °C and greater than or equal to -59 °C, then this field shall indicate the two's complement of the composite temperature of the NVM Subsystem in degrees Celsius. <del>Temperature measured in degrees Celsius is represented in two's complement (-1 °C to -59 °C)</del></td> </tr> </tbody> </table>	Value	Description	00h to 7Eh	<del>Temperature is measured in degrees Celsius (If the composite temperature of the NVM Subsystem is greater than or equal to 0 °C and less than or equal to 126 °C),</del> then this field shall indicate the composite temperature of the NVM Subsystem in degrees Celsius.	7Fh	If the composite temperature of the NVM Subsystem is greater than or equal to 127 °C <del>or higher,</del> then this field shall indicate a value of 7Fh.	80h	<del>No temperature data or temperature data</del> If the composite temperature of the NVM Subsystem is <del>more the</del> greater than 5 s old, then this field shall indicate a value of 80h.	81h	If the composite temperature of the NVM Subsystem is not accurate due to the failure of one or more temperature sensors, then this field shall indicate a value of 81h. <del>Temperature sensor failure</del>	82h to C3h	Reserved	C4h	If the composite temperature of the NVM Subsystem is less than or equal to <del>Temperature is</del> -60 °C <del>or lower,</del> then this field shall indicate a value of C4h.	C5h to FFh	If the composite temperature of the NVM Subsystem is less than or equal to -1 °C and greater than or equal to -59 °C, then this field shall indicate the two's complement of the composite temperature of the NVM Subsystem in degrees Celsius. <del>Temperature measured in degrees Celsius is represented in two's complement (-1 °C to -59 °C)</del>
Value	Description																
00h to 7Eh	<del>Temperature is measured in degrees Celsius (If the composite temperature of the NVM Subsystem is greater than or equal to 0 °C and less than or equal to 126 °C),</del> then this field shall indicate the composite temperature of the NVM Subsystem in degrees Celsius.																
7Fh	If the composite temperature of the NVM Subsystem is greater than or equal to 127 °C <del>or higher,</del> then this field shall indicate a value of 7Fh.																
80h	<del>No temperature data or temperature data</del> If the composite temperature of the NVM Subsystem is <del>more the</del> greater than 5 s old, then this field shall indicate a value of 80h.																
81h	If the composite temperature of the NVM Subsystem is not accurate due to the failure of one or more temperature sensors, then this field shall indicate a value of 81h. <del>Temperature sensor failure</del>																
82h to C3h	Reserved																
C4h	If the composite temperature of the NVM Subsystem is less than or equal to <del>Temperature is</del> -60 °C <del>or lower,</del> then this field shall indicate a value of C4h.																
C5h to FFh	If the composite temperature of the NVM Subsystem is less than or equal to -1 °C and greater than or equal to -59 °C, then this field shall indicate the two's complement of the composite temperature of the NVM Subsystem in degrees Celsius. <del>Temperature measured in degrees Celsius is represented in two's complement (-1 °C to -59 °C)</del>																
3	<p><b>Percentage Drive Life Used (PDLU):</b> <del>Contains</del> This field shall indicate an implementation-specific <del>a vendor-specific</del> 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, then the highest value <del>is shall be</del> 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 SMART / Health Information log page <del>NVMe SMART Health Log Page.</del></p>																

Figure 90: NVM Subsystem Health Data Structure (NSHDS)

Bytes	Description																																													
5:4	<p><b>Composite Controller Status (CCS):</b> This field shall indicate the Composite Controller Status Flags (refer to Figure TBD2) reports the composite status of all Controllers in the NVM Subsystem.</p> <p>The bits in this field are cleared after the NVM Subsystem Health Data Structure (refer to Figure 90) is returned in a Success Response associated with an 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'.</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Reset</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>15:13</td> <td>0</td> <td>Reserved</td> </tr> <tr> <td>12</td> <td>0</td> <td><del><b>Critical Warning (CWARN):</b> This bit is set to '1' when the Critical Warning bit in the Controller Health Status Changed Flags transitions from '0' to '1' in one or more Controllers in the NVM Subsystem.</del></td> </tr> <tr> <td>11</td> <td>0</td> <td><del><b>Available Spare (SPARE):</b> This bit is set to '1' when the Available Spare bit in the Controller Health Status Changed Flags transitions from '0' to '1' in one or more Controllers in the NVM Subsystem.</del></td> </tr> <tr> <td>10</td> <td>0</td> <td><del><b>Percentage Used (PDLU):</b> This bit is set to '1' when the Percentage Used field in the Controller Health Status Changed Flags transitions from '0' to '1' in one or more Controllers in the NVM Subsystem.</del></td> </tr> <tr> <td>09</td> <td>0</td> <td><del><b>Composite Temperature Change (CTEMP):</b> This bit is set to '1' when the Composite Temperature field in the Controller Health Status Changed Flags transitions from '0' to '1' in one or more Controllers in the NVM Subsystem.</del></td> </tr> <tr> <td>08</td> <td>Hwlnit</td> <td><del><b>Controller Status Change (CSTS):</b> This bit is set to '1' when the Controller Status field in the Controller Health Status Changed Flags transitions from '0' to '1' in one or more Controllers in the NVM Subsystem.</del></td> </tr> <tr> <td>07</td> <td>Hwlnit</td> <td><del><b>Firmware Activated (FA):</b> 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.</del></td> </tr> <tr> <td>06</td> <td>0</td> <td><del><b>Namespace Attribute Changed (NAC):</b> This bit is set to '1' when the 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.</del></td> </tr> <tr> <td>05</td> <td>0</td> <td><del><b>Controller Enable Change Occurred (CECO):</b> 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.</del></td> </tr> <tr> <td>04</td> <td>Hwlnit</td> <td><del><b>NVM Subsystem Reset Occurred (NSSRO):</b> 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.</del></td> </tr> <tr> <td>03</td> <td>0</td> <td>Reserved</td> </tr> <tr> <td>02</td> <td>0</td> <td><del><b>Shutdown Status (SHST):</b> 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.</del></td> </tr> <tr> <td>01</td> <td>Hwlnit</td> <td><del><b>Controller Fatal Status (CFS):</b> 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.</del></td> </tr> <tr> <td>00</td> <td>0</td> <td><del><b>Ready (RDY):</b> 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.</del></td> </tr> </tbody> </table>	Bits	Reset	Description	15:13	0	Reserved	12	0	<del><b>Critical Warning (CWARN):</b> This bit is set to '1' when the Critical Warning bit in the Controller Health Status Changed Flags transitions from '0' to '1' in one or more Controllers in the NVM Subsystem.</del>	11	0	<del><b>Available Spare (SPARE):</b> This bit is set to '1' when the Available Spare bit in the Controller Health Status Changed Flags transitions from '0' to '1' in one or more Controllers in the NVM Subsystem.</del>	10	0	<del><b>Percentage Used (PDLU):</b> This bit is set to '1' when the Percentage Used field in the Controller Health Status Changed Flags transitions from '0' to '1' in one or more Controllers in the NVM Subsystem.</del>	09	0	<del><b>Composite Temperature Change (CTEMP):</b> This bit is set to '1' when the Composite Temperature field in the Controller Health Status Changed Flags transitions from '0' to '1' in one or more Controllers in the NVM Subsystem.</del>	08	Hwlnit	<del><b>Controller Status Change (CSTS):</b> This bit is set to '1' when the Controller Status field in the Controller Health Status Changed Flags transitions from '0' to '1' in one or more Controllers in the NVM Subsystem.</del>	07	Hwlnit	<del><b>Firmware Activated (FA):</b> 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.</del>	06	0	<del><b>Namespace Attribute Changed (NAC):</b> This bit is set to '1' when the 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.</del>	05	0	<del><b>Controller Enable Change Occurred (CECO):</b> 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.</del>	04	Hwlnit	<del><b>NVM Subsystem Reset Occurred (NSSRO):</b> 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.</del>	03	0	Reserved	02	0	<del><b>Shutdown Status (SHST):</b> 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.</del>	01	Hwlnit	<del><b>Controller Fatal Status (CFS):</b> 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.</del>	00	0	<del><b>Ready (RDY):</b> 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.</del>
	Bits	Reset	Description																																											
	15:13	0	Reserved																																											
	12	0	<del><b>Critical Warning (CWARN):</b> This bit is set to '1' when the Critical Warning bit in the Controller Health Status Changed Flags transitions from '0' to '1' in one or more Controllers in the NVM Subsystem.</del>																																											
	11	0	<del><b>Available Spare (SPARE):</b> This bit is set to '1' when the Available Spare bit in the Controller Health Status Changed Flags transitions from '0' to '1' in one or more Controllers in the NVM Subsystem.</del>																																											
	10	0	<del><b>Percentage Used (PDLU):</b> This bit is set to '1' when the Percentage Used field in the Controller Health Status Changed Flags transitions from '0' to '1' in one or more Controllers in the NVM Subsystem.</del>																																											
	09	0	<del><b>Composite Temperature Change (CTEMP):</b> This bit is set to '1' when the Composite Temperature field in the Controller Health Status Changed Flags transitions from '0' to '1' in one or more Controllers in the NVM Subsystem.</del>																																											
	08	Hwlnit	<del><b>Controller Status Change (CSTS):</b> This bit is set to '1' when the Controller Status field in the Controller Health Status Changed Flags transitions from '0' to '1' in one or more Controllers in the NVM Subsystem.</del>																																											
	07	Hwlnit	<del><b>Firmware Activated (FA):</b> 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.</del>																																											
	06	0	<del><b>Namespace Attribute Changed (NAC):</b> This bit is set to '1' when the 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.</del>																																											
	05	0	<del><b>Controller Enable Change Occurred (CECO):</b> 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.</del>																																											
	04	Hwlnit	<del><b>NVM Subsystem Reset Occurred (NSSRO):</b> 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.</del>																																											
	03	0	Reserved																																											
	02	0	<del><b>Shutdown Status (SHST):</b> 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.</del>																																											
01	Hwlnit	<del><b>Controller Fatal Status (CFS):</b> 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.</del>																																												
00	0	<del><b>Ready (RDY):</b> 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.</del>																																												
7:6	Reserved																																													

### 5.7 Read NVM-MI Data Structure

Modify Figure 94 (NVM Subsystem Information Data Structure) as follows:

**Figure 94: NVM Subsystem Information Data Structure**

Bytes	Description						
00	<b>Number of Ports (NUMP):</b> This field indicates the maximum number of ports of any type supported by the NVM Subsystem. This is a 0's based value. The value of FFh is not supported because a port identifier of 256 is not able to be reported (refer to section 5.1.1).						
01	<b>NVMe-MI Major Version Number (MJR):</b> This field shall be set to 1h to indicate the major version number of this specification.						
02	<b>NVMe-MI Minor Version Number (MNR):</b> This field shall be set to 2h to indicate the minor version number of this specification.						
03	<p><b>NVMe-MI NVM Subsystem Capabilities (NNSC):</b> This field indicates the NVMe-MI capabilities of the NVM Subsystem.</p> <table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>7:1</td> <td>Reserved</td> </tr> <tr> <td>0</td> <td> <p><b>Status Reporting Enhancements (SRE):</b> If the status reporting enhancements are supported, then this bit shall be set to '1'. If the status reporting enhancements are not supported, then this bit is reserved.</p> <p>Implementations compliant to versions of this specification later than 1.2 shall set this bit to '1'.</p> <p>The status reporting enhancements consist of the following:</p> <ul style="list-style-type: none"> <li>• specified in the Get State Control Primitive that the NVM Subsystem Reset Occurred status bit is only set if a new firmware image is not activated to align to the behavior of NVM Subsystem Reset Occurred in the rest of this specification and the NVM Express Base Specification;</li> <li>• specified that a Controller Health Status Poll command with the Clear Changed Flags bit set to '1' also clears the Namespace Attribute Changed and Firmware Activated bits in Controller Health Data Structure to '0';</li> <li>• changed the Controller Enable Change Occurred bit in the Controller Health Data Structure from a status bit that is set when the Controller is enabled or disabled to a state bit that indicates if the Controller is currently enabled or disabled; and</li> <li>• set the:                             <ul style="list-style-type: none"> <li>○ Controller Enable Change Occurred bit;</li> <li>○ Critical Warning bit;</li> <li>○ Controller Status Change bit; or</li> <li>○ Controller Fatal Status bit;</li> </ul>                             in the Controller Health Status Changed Flags field whenever the corresponding bit in the Controller Health Data Structure changes state from '0' to '1' or '1' to '0' instead of only when it transitions from '0' to '1'.                         </li> </ul> </td> </tr> </tbody> </table>	Bits	Description	7:1	Reserved	0	<p><b>Status Reporting Enhancements (SRE):</b> If the status reporting enhancements are supported, then this bit shall be set to '1'. If the status reporting enhancements are not supported, then this bit is reserved.</p> <p>Implementations compliant to versions of this specification later than 1.2 shall set this bit to '1'.</p> <p>The status reporting enhancements consist of the following:</p> <ul style="list-style-type: none"> <li>• specified in the Get State Control Primitive that the NVM Subsystem Reset Occurred status bit is only set if a new firmware image is not activated to align to the behavior of NVM Subsystem Reset Occurred in the rest of this specification and the NVM Express Base Specification;</li> <li>• specified that a Controller Health Status Poll command with the Clear Changed Flags bit set to '1' also clears the Namespace Attribute Changed and Firmware Activated bits in Controller Health Data Structure to '0';</li> <li>• changed the Controller Enable Change Occurred bit in the Controller Health Data Structure from a status bit that is set when the Controller is enabled or disabled to a state bit that indicates if the Controller is currently enabled or disabled; and</li> <li>• set the:                             <ul style="list-style-type: none"> <li>○ Controller Enable Change Occurred bit;</li> <li>○ Critical Warning bit;</li> <li>○ Controller Status Change bit; or</li> <li>○ Controller Fatal Status bit;</li> </ul>                             in the Controller Health Status Changed Flags field whenever the corresponding bit in the Controller Health Data Structure changes state from '0' to '1' or '1' to '0' instead of only when it transitions from '0' to '1'.                         </li> </ul>
	Bits	Description					
7:1	Reserved						
0	<p><b>Status Reporting Enhancements (SRE):</b> If the status reporting enhancements are supported, then this bit shall be set to '1'. If the status reporting enhancements are not supported, then this bit is reserved.</p> <p>Implementations compliant to versions of this specification later than 1.2 shall set this bit to '1'.</p> <p>The status reporting enhancements consist of the following:</p> <ul style="list-style-type: none"> <li>• specified in the Get State Control Primitive that the NVM Subsystem Reset Occurred status bit is only set if a new firmware image is not activated to align to the behavior of NVM Subsystem Reset Occurred in the rest of this specification and the NVM Express Base Specification;</li> <li>• specified that a Controller Health Status Poll command with the Clear Changed Flags bit set to '1' also clears the Namespace Attribute Changed and Firmware Activated bits in Controller Health Data Structure to '0';</li> <li>• changed the Controller Enable Change Occurred bit in the Controller Health Data Structure from a status bit that is set when the Controller is enabled or disabled to a state bit that indicates if the Controller is currently enabled or disabled; and</li> <li>• set the:                             <ul style="list-style-type: none"> <li>○ Controller Enable Change Occurred bit;</li> <li>○ Critical Warning bit;</li> <li>○ Controller Status Change bit; or</li> <li>○ Controller Fatal Status bit;</li> </ul>                             in the Controller Health Status Changed Flags field whenever the corresponding bit in the Controller Health Data Structure changes state from '0' to '1' or '1' to '0' instead of only when it transitions from '0' to '1'.                         </li> </ul>						
31:04	Reserved						

**Modify a portion of Appendix C as follows:**

## Appendix C Example NVMe-MI Messages over SMBus/I2C

...

**Example 3:** In this example, a Management Controller issues an NVM Subsystem Health Status Poll command and clears the Composite Controller Status. Note that the MCTP packet sequence number is incremented from the last packet the Management Controller sent in Example 1. The NVMe-MI Message Type value of 08h with Opcode 01h makes this an NVM Subsystem Health Status Poll command. Bit 31 of Dword1 set to '1' clears the Composite Controller Status **Flags field** after preparing the response. Only the first non SR-IOV PCI function with any of the trigger able changes is requested.

...

**Example 6:** This example shows an NVMe Storage Device sending an acknowledgement Response Message to the Replay Control Primitive and then sending a second Response Message that replays the previous Response Message from specified offset of 0h. Note that the previous command is not reissued because that could return different data after having the Composite Controller Status **Flags field** cleared.

### Description of Specification Changes to TP6027

#### 4.2.1.4 Get State

**Modify Figure 42 (Get State Control Primitive Success Response Fields) as follows:**

**Figure 42: Get State Control Primitive Success Response Fields**

Bytes	Description
07:06	<b>Control Primitive Specific Response (CPSR):</b> This field <del>is shall indicate the contents of the Management Endpoint State data structure (refer to Figure TBD1) is used to return Control Primitive specific stats used to return Management Endpoint status. A Management Endpoint Reset of the corresponding Management Endpoint shall clear this field to 0h.</del>
...	...

### 5.3 Controller Health Status Poll

**Modify a portion of Figure 80 as follows:**

**Figure 80: Controller Health Data Structure (CHDS)**

Description		
<b>Controller Status (CSTS):</b> This field reports the Controller status.		
Bits	Reset <sup>1</sup>	Description
...	...	...
05	0 HwInit	<b>Controller Enable Change Occurred (CECO):</b> This bit <del>is set to '1' when the Enable bit shall indicate the value of the Enable bit (refer to CC.EN in the NVM Express Base Specification) changes state.</del>  Note that the name of this bit does not match the functionality, but the original name of this bit has been retained for historical continuity. Refer to version 1.2 of this specification for the original definition of this bit.
...	...	...

### 5.6 NVM Subsystem Health Status Poll

**Modify a portion of Figure 90 as follows:**

**Figure 90: NVM Subsystem Health Data Structure (NSHDS)**

Bytes	Description	
0	<b>NVM Subsystem Status (NSS):</b> This field indicates the status of the NVM Subsystem.	
	Bits	Description
	...	...
	4	<p><del>Reset Not Required (RNR): If this bit is cleared to '0', then the NVM Subsystem has experienced an error that prevents continued normal operation, then this bit shall be set to '1'. An NVM Subsystem Reset is required to resume normal operation.</del></p> <p>If this bit is set to '1', then the NVM Subsystem does require an NVM Subsystem Reset to resume normal operation, then this bit shall be cleared to '0' has not experienced an error that prevents continued normal operation. An NVM Subsystem Reset is not required to resume normal operation.</p>
...	...	
...	...	
<p><b>NOTES:</b></p> <p><del>1. An NVM Subsystem Reset shall reset the instance of the Controller Status bits dedicated to the out-of-band mechanism and the instance of the Controller Status bits dedicated to each Controller in the NVM Subsystem.</del></p> <p><del>The instance of the Controller Status bits dedicated to a Controller shall be reset by a Controller Level Reset (refer to the NVM Express Base Specification) of that Controller. Note that a Controller Level Reset may affect Controller Status bits in the out-of-band mechanism (e.g., a Controller Level Reset causes the CECO bit in the instance of the Controller Status bits dedicated to the out-of-band mechanism to be set to '1').</del></p> <p><del>No instance of the Controller Status bits shall be reset by any other resets other than the resets documented by this note.</del></p>		

**Modify a portion of Section 8.3.2 as follows:**

### 8.3.2 Controller Level Reset

Additional requirements and recommendations for Controller Level Resets are specified elsewhere in this specification. For example, bits and fields that are dedicated to each Controller in the in-band tunneling mechanism are reset as defined in Figure 80, Figure 81, and ~~Figure 90~~ **Figure TBD2**.

**Modify a portion of Section 8.3.3 as follows:**

### 8.3.3 Management Endpoint Reset

Additional requirements and recommendations for Management Endpoint Resets are specified elsewhere in this specification. For example, a Management Endpoint Reset:

- resets bits and fields that are dedicated to the out-of-band mechanism as defined in Figure 80, Figure 81, and ~~Figure 90~~ **Figure TBD2**;
- resets the value of the MCTP Transmission Unit Size field as defined by Figure 69; and
- clears the Control Primitive Specific Response field to 0h as defined in Figure 42.

## Description of Specification Changes to TP6033

Modify Figure 42 in Section 4.2.1.4 as follows:

### 4.2.1.4 Get State

Figure 42: Get State Control Primitive Success Response Fields

Description		
<b>Control Primitive Specific Response (CPSR):</b> This field is used to return Control Primitive specific status.		
Bits	Command Slot Specific <sup>1</sup>	Description
15	No	<b>Pause Flag (PFLG):</b> <del>This bit indicates whether or not the Management Endpoint is paused.</del> If the Management Endpoint is paused, then this bit shall be set to '1'. If the Management Endpoint is not paused, then this bit shall be cleared to '0'.
...	...	...

## Description of Specification Changes for NVM Express Base Specification

Modify a portion of section 5.16.1 as follows:

### 5.16.1 Log Specific Information

...

Figure 202: Get Log Page – Log Page Identifiers

Log Identifier	Scope	Log Page Name	Reference Section
...			
06h	Controller <sup>3</sup>	Device Self-test <sup>5</sup>	5.16.1.7
	Domain / NVM subsystem <sup>4, 6</sup>		
07h	<del>Vendor Specific</del> Controller / NVM subsystem <sup>7</sup>	Telemetry Host-Initiated <sup>5</sup>	5.16.1.8
08h	<del>Vendor Specific</del> Controller / NVM subsystem <sup>7</sup>	Telemetry Controller-Initiated <sup>5</sup>	5.16.1.9
09h	Domain / NVM subsystem <sup>6</sup>	Endurance Group Information	5.16.1.10
...			
C0h to FFh	Vendor specific <sup>5</sup>		
<p><b>KEY:</b>  <del>Namespace – The log page contains information about a specific namespace.</del>  <del>Endurance Group – The log page contains information about a specific Endurance Group.</del>  <del>Controller – The log page contains information about the controller that is processing the command.</del>  <del>Domain – The log page contains information about the Domain.</del>  <del>NVM subsystem – The log page contains information about the NVM subsystem.</del>  <del>Vendor Specific – The log page contains information that is vendor specific.</del></p>			
<p><b>NOTES:</b>            1. For namespace identifiers of 0h or FFFFFFFFh.            2. For namespace identifiers other than 0h or FFFFFFFFh.            3. Bit 0 is cleared to '0' in the DSTO field in the Identify Controller data structure (refer to Figure 275).            4. Bit 0 is set to '1' in the DSTO field in the Identify Controller data structure.            5. Selection of a UUID may be supported. Refer to section 8.25.            6. For NVM subsystems that support multiple domains (refer to the MDS bit in the Identify Controller data structure, Figure 275), Domain scope information is returned.            7. The scope is defined in the header of the log page (refer to the Telemetry Host-Initiated Scope field and the Telemetry Controller-Initiated Scope field).</p>			

...

Modify a portion of section 5.16.1.8 as follows:

#### 5.16.1.8 Telemetry Host-Initiated (Log Identifier 07h)

This log consists of a header (i.e., bytes 511:0 of the log page) describing the log and zero or more Telemetry Data Blocks (refer to section 8.24). The header shall always be available even if there is no Telemetry Host-Initiated Data available. All Telemetry Data Blocks are 512 bytes in size. The controller shall initiate a capture of the controller's internal controller state or internal NVM subsystem state to this log if the controller processes a Get Log Page command for this log with the Create Telemetry Host-Initiated

Data bit set to '1' in the Log Specific field. If the host specifies a Log Page Offset Lower value that is not a multiple of 512 bytes in the Get Log Page command for this log, then the controller shall abort the command with a status code of Invalid Field in Command. ~~This log page is global to the controller or global to the NVM subsystem.~~

...

**Figure 215: Telemetry Host-Initiated Log Page**

Bytes	Description										
...	...										
<del>380</del> 379:20	Reserved										
380	<p><b>Telemetry Host-Initiated Scope:</b> This field shall indicate the scope of the Telemetry Host-Initiated log page. Implementations compliant with versions of this specification later than 2.0 shall not set this field to a value of 0h.</p> <table border="1"> <thead> <tr> <th>Value</th> <th>Scope</th> </tr> </thead> <tbody> <tr> <td>00h</td> <td>Not reported</td> </tr> <tr> <td>01h</td> <td>Controller</td> </tr> <tr> <td>02h</td> <td>NVM subsystem</td> </tr> <tr> <td>03h to FFh</td> <td>Reserved</td> </tr> </tbody> </table>	Value	Scope	00h	Not reported	01h	Controller	02h	NVM subsystem	03h to FFh	Reserved
Value	Scope										
00h	Not reported										
01h	Controller										
02h	NVM subsystem										
03h to FFh	Reserved										
381	<p><b>Telemetry Host-Initiated Data Generation Number:</b> Contains a value that is incremented each time the <del>controller captures its</del> internal controller state <del>or internal NVM subsystem state</del> for this log page is captured. If the value of this field is FFh, then the field shall be cleared to 0h when incremented (i.e., rolls over to 0h).</p>										
382	<p><b>Telemetry Controller-Initiated Data Available:</b> Contains the value of Telemetry Controller-Initiated Data Available field in the Telemetry Controller-Initiated log (refer to Figure 216).</p>										
383	<p><b>Telemetry Controller-Initiated Data Generation Number:</b> Contains the value of the Telemetry Controller-Initiated Data Generation Number field in the Telemetry Controller-Initiated log (refer to Figure 216).</p>										
...	...										

**Modify a portion of section 5.16.1.9 as follows:**

#### 5.16.1.9 Telemetry Controller-Initiated (Log Identifier 08h)

This log consists of a header (i.e., bytes 511:0 of the log page) describing the log and zero or more Telemetry Data Blocks (refer to section 8.23). ~~The header shall always be available even if there is no Telemetry Controller-Initiated Data available.~~ All Telemetry Data Blocks are 512 bytes in size. This log is a ~~controller-initiated~~ ~~controller-initiated~~ capture of the controller's internal state ~~or NVM subsystem's internal state~~. The Telemetry Controller-Initiated Data for Data Area 1 through Data Area 3 shall persist across all resets. The Telemetry Controller-Initiated Data for Data Area 4 may persist across controller resets. If the host specifies a Log Page Offset Lower value that is not a multiple of 512 bytes in the Get Log Page command for this log, then the controller shall return an error of Invalid Field in Command. ~~This log page is global to the controller.~~

**Modify a portion of Figure 216 as follows:**

**Figure 216: Get Log Page – Telemetry Controller-Initiated Log (Log Identifier 08h)**

Bytes	Description
...	...

**Figure 216: Get Log Page – Telemetry Controller-Initiated Log (Log Identifier 08h)**

Bytes	Description										
<del>381</del> 380:20	Reserved										
381	<p><b>Telemetry Controller-Initiated Scope:</b> This field shall indicate the scope of the Telemetry Controller-Initiated log page. Implementations compliant with versions of this specification later than 2.0 shall not set this field to a value of 0h.</p> <table border="1"> <thead> <tr> <th>Value</th> <th>Scope</th> </tr> </thead> <tbody> <tr> <td>00h</td> <td>Not reported</td> </tr> <tr> <td>01h</td> <td>Controller</td> </tr> <tr> <td>02h</td> <td>NVM subsystem</td> </tr> <tr> <td>03h to FFh</td> <td>Reserved</td> </tr> </tbody> </table>	Value	Scope	00h	Not reported	01h	Controller	02h	NVM subsystem	03h to FFh	Reserved
Value	Scope										
00h	Not reported										
01h	Controller										
02h	NVM subsystem										
03h to FFh	Reserved										
382	<p><b>Telemetry Controller-Initiated Data Available:</b> If this field is cleared to 0h, then the log page consists of only the 512-byte header and does not contain saved internal controller state or saved internal NVM subsystem state available to be reported in response to a Get Log Page command issued to the Admin Submission Queue.</p> <p>If this field is set to 1h, then the log page contains saved internal controller state or saved internal NVM subsystem state to be reported in response to a Get Log Page command issued to the Admin Submission Queue. If this field is set to 1h, then it shall not be cleared to 0h until a Get Log Page command with Retain Asynchronous Event bit cleared to '0' for the Telemetry Controller-Initiated log page completes successfully. This value is persistent across power states and reset.</p> <p>Regardless of the value of this field, the log page may contain saved internal controller state or saved internal NVM subsystem state available to be reported in response to a Get Log Page command issued out-of-band to a Management Endpoint (refer to the NVM Express Management Interface Specification).</p> <p>Other values are reserved.</p>										
383	<p><b>Telemetry Controller-Initiated Data Generation Number:</b> Contains a value that is incremented each time the <del>controller initiates a capture of its</del> internal controller state or the NVM subsystem state for this log page is captured <del>into the Telemetry Controller-Initiated Data Blocks</del>. If the value of this field is FFh, then the field shall be cleared to 0h when incremented (i.e., rolls over to 0h). This field is persistent across power cycles.</p>										
...	...										

## 8.24 Telemetry

**Modify a portion of section 8.24 as follows:**

...

The preparation, collection, and submission of telemetry data is similar for host-initiated and controller-initiated data; the primary difference is the trigger for the collection. The operational model for telemetry is:

1. The host identifies controller support for Telemetry log pages in the Identify Controller data structure;
2. The host may indicate the support for the Telemetry Host-Initiated Data Area 4 and Telemetry Controller-Initiated Data Area 4 by setting the Extended Telemetry Data Area 4 Supported (ETDAS) field to 1h in the Host Behavior Support feature (refer to section 5.27.1.18);
3. The host prepares an area to store telemetry data if needed;
4. To receive notification that controller-initiated telemetry data is available, the host enables Telemetry Log Notices using the Asynchronous Event Configuration feature (refer to section 5.27.1.8); and
5. If the host decides to collect host-initiated telemetry data or the controller signals that controller-initiated telemetry data is available:

- a. The host reads the appropriate blocks of the Telemetry Data Area from the host-initiated log (refer to section 5.16.1.8) or the controller-initiated log (refer to section 5.16.1.9). If possible, the host should collect Telemetry Data Area 1, 2, 3, and 4. The host reads the log in 512 byte Telemetry Data Block units. The host should set the Retain Asynchronous Event bit to '1';
- b. The host re-reads the header of the log page and ensures that the Telemetry Host-Initiated Data Generation Number field from the host-initiated log or the Telemetry Controller-Initiated Data Generation Number field in the controller-initiated log matches the original value read. If these values do not match, then the data captured is not consistent and should be re-read from the log page with the Retain Asynchronous Event bit set to '1';
- c. If the host is reading the controller-initiated log, then the host ensures that the Telemetry Controller-Initiated Data Available field is still set to 1h after reading the appropriate blocks of the Telemetry Data Area because the Telemetry Controller-Initiated Data Available field may have been cleared to 0h by another entity while the log page was being read;
- d. ~~e.~~ If the host is reading the controller-initiated log, then the host reads any portion of that log page with the Retain Asynchronous Event bit cleared to '0' to indicate to the controller that the host has completed reading the controller-initiated log page; and
- e. ~~d.~~ When all telemetry data has been saved, the data should be forwarded to the manufacturer of the controller.

The trigger for the collection for host-initiated data is typically a system crash, but may also be initiated during normal operation. The host proceeds with a host-initiated data collection by submitting the Get Log Page command for the Telemetry Host-Initiated log page with the Create Telemetry Host-Initiated Data bit set to '1' in the Log Specific field. The controller should complete the command quickly (e.g., in less than one second) to avoid a user rebooting the system prior to completion of the data collection.

The NVM subsystem is allowed to provide a host-initiated log page per controller or a shared host-initiated log page across all controllers in the NVM subsystem. When a shared host-initiated log is implemented, the Telemetry Host-Initiated Data Generation Number field in the host-initiated log is used to allow the host to detect that the Telemetry Host-Initiated log has been changed by a host from a different controller.

The controller notifies the host to collect controller-initiated data through the completion of an Asynchronous Event Request command with an Asynchronous Event Type of Notice that indicates a Telemetry Log Changed event. The host may also determine controller-initiated data is available via the Telemetry Controller-Initiated Data Available field in the Telemetry Host-Initiated or the Telemetry Controller-Initiated log pages. The host proceeds with a controller-initiated data collection by submitting the Get Log Page command for the Telemetry Controller-Initiated log page. Once the host has started reading the Telemetry Controller-Initiated log page, the controller should avoid modifying the controller-initiated data until the host has finished reading all controller-initiated data. The amount of time for the host to read the controller-initiated data is vendor specific.

Since there is only one set of controller-initiated data, the controller is responsible for prioritizing the version of the controller-initiated data that is available for the host to collect. When the controller replaces the controller-initiated data with new controller-initiated data, the controller shall increment the Telemetry Controller-Initiated Data Generation Number field. The host needs to ensure that the Telemetry Controller-Initiated Data Generation Number field has not changed between the start and completion of the controller-initiated data collection to ensure the data captured is consistent.