



LEGAL NOTICE:

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

This NVM Express revision 1.3 technical proposal is proprietary to the NVM Express, Inc. (also referred to as "Company") and/or its successors and assigns.

NOTICE TO USERS WHO ARE NVM EXPRESS, INC. MEMBERS: Members of NVM Express, Inc. have the right to use and implement this NVM Express revision 1.3 technical proposal subject, however, to the Member's continued compliance with the Company's Intellectual Property Policy and Bylaws and the Member's Participation Agreement.

NOTICE TO NON-MEMBERS OF NVM EXPRESS, INC.: If you are not a Member of NVM Express, Inc. and you have obtained a copy of this document, you only have a right to review this document or make reference to or cite this document. Any such references or citations to this document must acknowledge NVM Express, Inc. copyright ownership of this document. The proper copyright citation or reference is as follows: "© 2007 - 2019 NVM Express, Inc. ALL RIGHTS RESERVED." When making any such citations or references to this document you are not permitted to revise, alter, modify, make any derivatives of, or otherwise amend the referenced portion of this document in any way without the prior express written permission of NVM Express, Inc. Nothing contained in this document shall be deemed as granting you any kind of license to implement or use this document or the specification described therein, or any of its contents, either expressly or impliedly, or to any intellectual property owned or controlled by NVM Express, Inc., including, without limitation, any trademarks of NVM Express, Inc.

LEGAL DISCLAIMER:

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

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

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

NVM Express Technical Proposal for New Feature

Technical Proposal ID	4042a – Further Persistent Event Log Events
Change Date	29 April 2019
Builds on Specification	NVM Express 1.3c
References Ratified TPs	TP 4004 , ANA Base Protocol TP 4007 , Persistent Event Log TP 4016 , Rebuild Assist TP 4018a , NVM Sets and Read Recovery Level TP 4027 , UUIDs for VS Information TP 4033 , Enhanced Command Retry TP 4050 , Endurance Group Info Enhancements TP 4039 , Administrative Controller

Technical Proposal Author(s)

Name	Company
David Derosa, Curtis Ballard, Olga Otte	Hewlett Packard Enterprise
David Black, Austin Bolen, Gerry Fredette	Dell EMC
Jim Hatfield	Seagate
Paul Suhler	Micron
Michael Allison	Intel

This technical proposal defines additional events that are able to be reported in the Persistent Event Log defined in TP 4007.

The following new persistent event log events are described in this TP:

- Subsystem Hardware Error;
- Set Feature;
- Format;
- Sanitize;
- Namespace Create/Delete;
- Telemetry log created;
- TCG; and
- Temperature Excursion.

Editorial Conventions: Black text indicates existing text or entirely new content if designated as new content with **dark green comments** at start of the new content. New content mixed with existing content is shown in **blue text**. Existing text to be removed is shown in ~~red strikethrough text~~. Text to be moved is shown in **brown text**. Comments intended only for discussion or explanation and not for inclusion in the specification are shown in **dark green text**. Items to be determined during integration of the technical proposal into the specification, e.g. figure numbering, are shown in **red text**.

Revision History

Revision Date	Change Description
2/21/2018	First version
2/28/2018	Converted NVM Subsystem Hardware Error skeleton into draft event format Rough draft of Change Namespace Event added
3/7/2018	Renamed to TP 4042 instead of TP 4007b Rough draft of Format Event incorporated Nearly complete draft of NVM Subsystem Hardware Error event
4/4/2018	Updated NVM Subsystem Hardware Error Event with additional PCIe conditions and struck out potential hardware error codes that were determined to be not appropriate for this log.
4/4/2018 - Clean	Same document but closed comments/strikeouts removed for easier reading
4/18/2018	Reworked the Change Namespace event following 4-18 team discussion Made the editorial changes that were requested in TP 4007
4/25/2018	Accepted changes to Change Namespace event Corrections to Change Namespace event from Gerry Fredette
5/17/2018	New Set Feature event and support reporting method from Mike Allison
5/25/2018	Cleaned up changes that had been discussed in small group meeting Added a TCG event
6/21/2018	Removed list of supported features for Set Feature event Added a bit to Set Feature to report response data format Editorial changes in the Set Feature descriptive text Merged in updates to Format event from David Derosa Make AER information reporting in PCIe error events options
6/27/2018	Split Format event into two events for format operation start and format operation complete Moved TCG defined event after vendor defined event Revised PCIe link change and link not active PCI Error types to clarify that they are only logged for exceptions, not normal transitions
7/10/2018	Refined the Format Event Added the Sanitize Event
7/13/2018	Cleaned up reviewed changes Incorporated review comments from Mike Allison
7/13b/2018	Discussed/resolved remaining open comments from Mike Allison Reverted Format and Sanitize events back to only being logged when an operation occurred and flagged for discussion in full technical WG meeting First pass of proposed optional/mandatory requirements for these new events
8/17/2018	Draft of potential method for reporting which optional features inside of events that can be reported are supported (e.g., which hardware errors are able to be detected)
8/29/2018	Removed optional feature reported method from 8/17 draft – small group decided it was too much complexity for the value. Cleaned up resolved comments First pass of allowed/prohibited features for Set Features logging 8/29+ Added clarifying statement for hardware error events that what hardware errors are able to be detected is vendor specific
9/4/2018	Added Telemetry Log Create Event from Gerry Fredette, Dell EMC Incorporated review comments from Mike Allison, Intel Reworked rules for Format NVM Start Event creation Adjusted references to core spec revision 1.3c
9/10/2018	Completed Set Feature log recommendations table Moved Set Features log recommendations into existing features tables Reviewed and made minor revisions to Telemetry Log Created event
9/13/2018	Pulled in Vendor Specific Event from TP 4007 to consider UUID index Several clarifications from Fred Knight review Removed option of saving Get Features data instead of Set Features data

10/8/2018	Accepted reviewed changed Merged in proposed Thermal Excursion event from WD
11/19/2018	Included alternate Thermal Excursion response data format from WD New text in Vendor Specific event for how to handle a UUID Index
11/28/2018	Incorporated review comments from Mike Allison, Intel
12/5/2018	Clean version for phase 2 exit closing all open tech questions <ul style="list-style-type: none"> • Thermal Excursion event – chose bitmap format response data • No change to optional/mandatory table • Endurance group critical warning text – no change • Set Feature event report completion DWord 0 – no change • Removed UUID Index from Vendor Specific Event
12/12/2018	Cleaned up change tracking for phase 2 exit vote Edited Thermal Excursion event to use an enumerated list instead of bitmap
1/18/2019	Editorial changes in preparation for 30 day member review Removed editorial comments from phase 2 discussions
2/27/2019	Added integration definition with TP 4039 Administrating Controller.
2/27/2019	Integration
3/2/2019	Review/acceptance of incorporation changes Editorial changes from 2/28 technical WG call
3/13/2019	Fixing byte offsets for the Thermal Excursion Event format.
4/10/2019	Correct color coding of text to match definition. Updated color coding definition to include red text as items to be defined during integration.
4/18/2019	Removed comments and accepted changed for ratification.
4/29/2019	Ratified

Description of Specification Changes

5.14.1 Log Specific Information

Figure 90 and Figure 91 define the Log pages that may be retrieved with the Get Log Page command.

Note to Editor: No changes until figure 90, unchanged intro text not included for easier reading

Note to Editor: Footnote 5 reference needs added to Persistent Event Log in the log page identifier table.

Figure 90: Get Log Page – Log Page Identifiers

Log Identifier	O/M	Scope	Description	Reference Section
00h	Reserved			
01h	M	Controller	Error Information	5.14.1.1
02h	M	NVM subsystem ¹	SMART / Health Information	5.14.1.2
	O	Namespace ²		
03h	M	NVM subsystem	Firmware Slot Information	5.14.1.3
04h	O	Controller	Changed Namespace List	5.14.1.4
05h	O	Controller	Commands Supported and Effects	5.14.1.5
06h	O	NVM subsystem	Device Self-test	5.14.1.6
07h	O	Controller	Telemetry Host-Initiated	5.14.1.7
08h	O	Controller	Telemetry Controller-Initiated	5.14.1.8
09h	O	NVM subsystem	Endurance Group Information	5.14.1.9
0Ah	O	NVM subsystem	Predictable latency Per NVM Set	5.14.1.10
0Bh	O	NVM subsystem	Predictable Latency Event Aggregate	5.14.1.11
0Ch	O	Controller	Asymmetric Namespace Access	5.14.1.12
0Dh	O	NVM subsystem	Persistent Event Log i.a.i.5	5.14.1.13
0Eh – 6Fh	Reserved			
70h	Discovery (refer to the NVMe over Fabrics specification)			
71h – 7Fh	Reserved for NVMe over Fabrics			
80h – BFh	I/O Command Set Specific			
C0h – FFh	Vendor specific			
KEY: O = Optional, M = Mandatory Namespace = The log page contains information about a specific namespace. Controller = The log page contains information about the controller that is processing the command. NVM subsystem = The log page contains information about the NVM subsystem.				
NOTES: 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 Error! Reference source not found.). 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 Error! Reference source not found.				

5.14.1.TBD.1 Persistent Event Log Events

The values that may be reported in the Event Type field (refer to section 5.14.1.TBD) of the event header for events in the Persistent Event log are defined in Figure TP4007 TBD + 4.

Figure TP4007 TBD + 4: Persistent Event Log Event Types

Value	O/M ¹	Definition	Reference Section
00h		Reserved	
01h	NOTE 2	SMART/Health Log Snapshot	5.14.1.TBD.1.1
02h	M	Firmware Commit	5.14.1.TBD.1.2
03h	M	Timestamp Change	5.14.1.TBD.1.3
04h	M	Power-on or Reset	5.14.1.TBD.1.4
05h	M	NVM Subsystem Hardware Error	5.14.1.TBD.1.5
06h	NOTE 3	Change Namespace	5.14.1.TBD.1.6
07h	NOTE 3	Format NVM Start	5.14.1.TBD.1.7
08h	NOTE 3	Format NVM Completion	5.14.1.TBD.1.8
09h	NOTE 3	Sanitize Start	5.14.1.TBD.1.9
0Ah	NOTE 3	Sanitize Completion	5.14.1.TBD.1.10
0Bh	O	Set Feature	5.14.1.TBD.1.11
0Ch	O	Telemetry Log Created	5.14.1.TBD.1.12
0Dh	O	Thermal Excursion	5.14.1.TBD.1.13
05h 0Eh to DDh		Reserved	
DEh	O	Vendor Specific Event	5.14.1.9.1.5 5.14.1.TBD.1.14
DFh	O	TCG Defined	5.14.1.TBD.1.15
DFh E0h to FFh		Reserved	
NOTES: 1. O/M definition: O = Optional, M = Mandatory 2. Mandatory for NVMe over PCIe, Optional for NVMe over Fabrics 3. Mandatory if the command used to initiate the activity reported by the event is supported			

Text not changed is not shown for easier reading

NOTE TO EDITOR: All of the following section 5.14 content is new text. It is shown in black for easier reading.

5.14.1.TBD.1.5 NVM Subsystem Hardware Error Event (Event Type 05h)

An NVM Subsystem Hardware Error event shall be recorded in the Persistent Event Log when a supported NVM subsystem hardware error event is detected. Which of the NVM subsystem hardware error events are supported is vendor specific. The NVM subsystem hardware error event shall set the Persistent Event Log Event Format Header:

- Event Type field to 05h; and
- Event Type Revision Field to 01h.

All detected NVM Subsystem Hardware Error events supported by the NVM Subsystem shall be logged unless otherwise specified (e.g., suppressed due to reoccurrence frequency (refer to section 5.14.1.TBD)). NVM Subsystem Hardware Error event fields reporting information that is not available (e.g., due to a PCIe optional feature that is not implemented) shall be set to zero unless otherwise specified in the NVM Subsystem Hardware Error Event code description.

The NVM Subsystem Hardware Error Event data is specified in Figure TBD + 10.

Figure TBD + 10: NVM Subsystem Hardware Error Event Format (Event Type 05h)

Bytes	Description
1:0	NVM Subsystem Hardware Error Event Code: This field contains a code (refer to Figure TBD + 11) indicating the type of NVM subsystem hardware error that is being reported.
3:2	Reserved
M+3:4	Additional Hardware Error Information: This field contains additional information about the hardware error event indicated in the NVM Subsystem Hardware Error Event Code field (refer to Figure TBD + 11). Where M is the number of bytes of additional hardware error information. This field is omitted if the subsystem hardware error being reported does not contain additional hardware error information.

Figure TBD + 11: NVM Subsystem Hardware Error Event Codes

Code	Description
00h	Reserved
01h	PCIe Correctable Error: Indicates that the NVM Subsystem has detected that a PCIe correctable error occurred. Refer to Figure TBD + 12 for the format of the Additional Hardware Error Information field.
02h	PCIe Uncorrectable Non-fatal Error: Indicates that the NVM Subsystem has detected that a PCIe uncorrectable non-fatal error occurred. Refer to Figure TBD + 12 for the format of the Additional Hardware Error Information field.

Code	Description								
03h	<p>PCIe Uncorrectable Fatal Error: Indicates that the NVM Subsystem has detected that a PCIe uncorrectable fatal error occurred.</p> <p>Refer to</p> <p>Figure TBD + 12 for the format of the Additional Hardware Error Information field.</p>								
04h	<p>PCIe Link Status Change: Indicates that a change in the values reported in the PCI Express Link Status register (refer to section 2.5.8) have changed due to an attempt to correct unreliable link operation.</p> <p>The Additional Hardware Error Information field shall be set to the contents of the PCI Express Link Status register at the time of the event.</p>								
05h	<p>PCIe Link Not Active: Indicates that the Data Link Control and Management State Machine (refer to the PCI Express Base specification) has transitioned out of the DL_Active state without a corresponding event (e.g., without an indication from the host that the link is to be disabled).</p> <p>This NVM subsystem hardware error event does not contain additional hardware error information.</p>								
06h	<p>Critical Warning Condition: Indicates the NVM subsystem has detected a condition that causes a bit in the Critical Warning field of the SMART / Health Information log (refer to section 5.14.1.2) to be set to one.</p> <p>Bits in this field represent the associated state at the time of this event.</p> <p>The Additional Hardware Error Information field shall be set at the time of the event using the same format as is specified for the Critical Warning field of the SMART / Health Information log.</p>								
07h	<p>Endurance Group Critical Warning Condition: Indicates that the NVM subsystem has detected a condition that causes a bit in the Critical Warning field of an Endurance Group Information log (refer to section TBD) to be set to one.</p> <p>Bits in this field represent the associated state at the time of this event.</p> <p>The Additional Hardware Error Information field shall be four bytes long.</p> <table border="1"> <thead> <tr> <th>Bytes</th><th>Definition</th></tr> </thead> <tbody> <tr> <td>0</td><td>Shall be set at the time of the event using the same format as is specified for the Critical Warning field of the Endurance Group Information log page.</td></tr> <tr> <td>1</td><td>Reserved</td></tr> <tr> <td>3:2</td><td>Shall be set to the Endurance Group Identifier for the associated endurance group.</td></tr> </tbody> </table> <p>Note to Editor: The Endurance group critical warning was introduced in TP 4050</p>	Bytes	Definition	0	Shall be set at the time of the event using the same format as is specified for the Critical Warning field of the Endurance Group Information log page.	1	Reserved	3:2	Shall be set to the Endurance Group Identifier for the associated endurance group.
Bytes	Definition								
0	Shall be set at the time of the event using the same format as is specified for the Critical Warning field of the Endurance Group Information log page.								
1	Reserved								
3:2	Shall be set to the Endurance Group Identifier for the associated endurance group.								
08h	<p>Unsafe Shutdown: Indicates that the controller incremented the Unsafe Shutdowns field value in the SMART / Health Information Log.</p> <p>The Additional Hardware Error Information field shall be set to the value from the Unsafe Shutdowns field in the SMART / Health Information log at the time of the event.</p>								
09h	<p>Controller Fatal Status: Indicates that the Controller Fatal Status (CSTS.CFS) bit has been set to '1'.</p> <p>This NVM subsystem hardware error event does not contain additional hardware error information.</p>								
0Ah	<p>Media and Data Integrity Status: Indicates that a completion queue entry contained a Media and Data Integrity status code (refer to Figure 36) other than 86h (i.e., Access Denied) or 87h (i.e., Deallocated or Unwritten logical Block).</p> <p>The Additional Hardware Error Information field shall be set to the contents of the completion queue entry.</p>								
0Bh to FFh	Reserved								

Figure TBD + 12: Additional Hardware Error Information for correctable and uncorrectable PCIe errors

Bytes	Description
01:00	PCIe Device Status Register: Contains the contents of the PCI Device Status Register (Refer to PCI Express specification) at the time of the event.
02	Bits 7:1 Reserved Bit 0 PCIe AER Supported: set to '1' indicates that PCIe AER (Refer to PCI Express Specification) is supported and that the PCIe AER Error Status field, PCIe AER Error Mask field, PCIe AER Header Log Register field, and the PCIe AER TLP Prefix Log Register field is reported. Bit 0 cleared to '0' indicates that PCIe AER is not supported and that the PCIe AER Error Status field, PCIe AER Error Mask field, PCIe AER Header Log Register field, and PCIe AER TLP Prefix Log Register field is not reported (i.e., bytes 80:16 are not reported).
15:03	Reserved
31:16	PCIe AER Error Status: Contains the contents of: a) the PCIe AER Correctable Error Status Register (refer to section 2.6.5) at the time of the event if the error is a correctable error; or b) The PCIe AER Uncorrectable Error Status Register (refer to section 2.6.2), at the time of the event if the error is an uncorrectable error.
47:32	PCIe AER Error Mask: Contains the contents of a) the PCIe AER Correctable Error Mask Register (refer to section 2.6.6) at the time of the event if the error is a correctable error; or b) the PCIe AER Uncorrectable Error Mask Register (refer to section 2.6.3) at the time of the event if the error is a uncorrectable error.
63:48	PCIe AER Header Log Register: Contains the contents of the PCIe AER Header Log Register (refer to section 2.6.8), if supported, at the time of the event.
80:64	PCIe AER TLP Prefix Log Register: Contains the contents of the PCIe AER TLP Prefix Log Register (refer to section 2.6.9), if supported, at the time of the event.

5.14.1.TBD.1.6 Change Namespace Event (Event Type 06h)

The Changed Namespace Event (refer to

Figure TBD + 13) persists the host parameters used for successful namespace management commands. The event contains a Persistent Event Log Event header and the Change Namespace Event data.

The Changed Namespace Event shall set the Persistent Event Log Event Format Header:

- Event Type field to 06h; and
- Event Type Revision Field to 01h.

Figure TBD + 13: Change Namespace Event Data Format (Event Type 06h)

Bytes	Description
03:00	Namespace Management CDW10: Contains the value from command Dword 10 of the Namespace Management command that initiated the namespace change event (refer to Figure 122).
07:04	Reserved
15:08	Namespace Size (NSZE): For a create operation, contains the NSZE value from the Identify Namespace data structure in the Namespace Management command (refer to Figure 123). For a delete operation that specifies a single namespace this field contains the value from the NSZE field of the Identify Namespace data (refer to Figure 109) for the namespace being deleted. For a delete operation that specifies all namespaces this field is reserved.
31:16	Namespace Capacity (NCAP): For a creation operation, contains the NCAP value from the Identify Namespace data structure in the Namespace Management command (refer to Figure123). For a delete operation that specifies a single namespace this field contains the value from the NCAP field of the Identify Namespace data (refer to Figure 109) for the namespace being deleted. For a delete operation that specifies all namespaces this field is reserved.
32	Formatted LBA Size (FLBAS): For a create operation, contains the FLBAS value from the Identify Namespace data structure in the Namespace Management command (refer to Figure123). For a delete operation that specifies a single namespace this field contains the value from the FLBAS field of the Identify Namespace data (refer to Figure 109) for the namespace being deleted. For a delete operation that specifies all namespaces this field is reserved.
33	End-to-end Data Protection Type Settings (DPS): For a create operation, contains the DPS value from the Identify Namespace data structure in the Namespace Management command (refer to Figure123). For a delete operation that specifies a single namespace this field contains the value from the DPS field of the Identify Namespace data (refer to Figure 109) for the namespace being deleted. For a delete operation that specifies all namespaces this field is reserved.
34	Namespace Multi-path I/O and Namespace Sharing Capabilities (NMIC): For a create operation, contains the NMIC value from the Identify Namespace data structure in the Namespace Management command (refer to Figure123). For a delete operation that specifies a single namespace this field contains the value from the NMIC field of the Identify Namespace data (refer to Figure 109) for the namespace being deleted. For a delete operation that specifies all namespaces this field is reserved.
35	Reserved
39:36	ANA Group Identifier (ANAGRPID): For a create operation, contains the ANAGRPID value from the Identify Namespace data structure in the Namespace Management command (refer to Figure 123), if specified, or contains the ANAGRPID value from the Identify Namespace data (Refer to Figure 109) after the namespace was created if an ANA Group Identifier was not specified in the command. For a delete operation that specifies a single namespace this field contains the value from the ANAGRPID field of the Identify Namespace data (refer to Figure 109) for the namespace being deleted. For a delete operation that specifies all namespaces this field is reserved. If ANA Groups are not supported, then the ANAGRPID field shall be cleared to 0h. <i>Editorial note: TP 4004, ANA Base Protocol, added the ANA Group Identifier field to the Identify Namespace data.</i>
41:40	NVM Set Identifier (NVMSETID): For a create operation, contains the NVMSETID value from the Identify Namespace data structure in the Namespace Management command (refer to Figure123). For a delete operation that specifies a single namespace this field contains the value from the NVMSETID field of the Identify Namespace data (refer to Figure 109) for the namespace being deleted. For a delete operation that specifies all namespaces this field is reserved. <i>Editorial note: TP 4018a, NVM Sets and Read Recovery Level, added the NVM Set Identifier field to the Identify Namespace data.</i>
43:42	Reserved
47:44	Namespace ID (NSID): For a create operation, contains the NSID for the namespace that was created. For a delete operation, contains the NSID from CDW1.NSID (i.e., the NSID for the namespace being deleted or FFFFFFFFh for a delete operation specifying all namespaces).

5.14.1.TBD.1.7 Format NVM Start Event (Event Type 07h)

A Format NVM Start event shall be recorded in the Persistent Event Log after successfully validating the command parameters of a Format NVM Command (refer to section 5.23) and before modifying any of the contents of the NVM.

The Format NVM Start event shall set the Persistent Event Log Event Format Header:

- Event Type field to 07h and
- Event Type Revision field to 01h

Figure TBD + 14: Format NVM Start Event Data Format (Event Type 07h)

Bytes	Description
03:00	Namespace Identifier: Contains the namespace identifier specified in the Format NVM command.
04	Format NVM Attributes (FNA): Contains the value from the identify controller FNA field.
07:04	Reserved
11:08	Format NVM CDW10: Contains the value from command Dword 10 of the Format NVM command (refer to Figure 173).

5.14.1.TBD.1.8 Format NVM Completion Event (Event Type 08h)

A Format NVM Completion event shall be recorded in the Persistent Event Log at the completion of a Format NVM command that resulted in modification of the contents of the NVM.

The Format NVM Completion event shall set the Persistent Event Log Event Format Header:

- Event Type field to 08h and
- Event Type Revision field to 01h

Figure TBD + 15: Format NVM Completion Event Data Format (Event Type 08h)

Bytes	Description								
03:00	Namespace Identifier: Contains the namespace identifier specified in the Format command.								
04	Smallest Format Progress Indicator: For a Format NVM command that formats a single namespace this field contains the lowest numerical value that was available for reporting in the FPI field of the Identify Namespace data structure (i.e. if the format did not complete successfully and the FPI field is supported then this field contains the percentage of the namespace that remained to be formatted at the time the format NVM command completed, refer to Figure 109) during the format operation. For a Format NVM command that formats all namespaces this field shall be set to 00h.								
05	Format NVM Status: Contains the status of the format operation. <table><tr><th>Bits</th><th>Definition</th></tr><tr><td>7:2</td><td>Reserved</td></tr><tr><td>1</td><td>Incomplete Format: If set to '1', then the format operation modified one or more LBAs but did not complete successfully. If set to '1', then the Format NVM Error bit shall be set to '1'. If cleared to '0' then the format operation either did not modify any LBAs or the format operation completed successfully.</td></tr><tr><td>0</td><td>Format NVM Error: If set to '1', then the format operation did not complete successfully. If cleared to '0', then the format operation completed successfully.</td></tr></table>	Bits	Definition	7:2	Reserved	1	Incomplete Format: If set to '1', then the format operation modified one or more LBAs but did not complete successfully. If set to '1', then the Format NVM Error bit shall be set to '1'. If cleared to '0' then the format operation either did not modify any LBAs or the format operation completed successfully.	0	Format NVM Error: If set to '1', then the format operation did not complete successfully. If cleared to '0', then the format operation completed successfully.
Bits	Definition								
7:2	Reserved								
1	Incomplete Format: If set to '1', then the format operation modified one or more LBAs but did not complete successfully. If set to '1', then the Format NVM Error bit shall be set to '1'. If cleared to '0' then the format operation either did not modify any LBAs or the format operation completed successfully.								
0	Format NVM Error: If set to '1', then the format operation did not complete successfully. If cleared to '0', then the format operation completed successfully.								
07:06	Completion Information: Contains a vendor specific value that may provide more information about the completion of the format operation (e.g., if the sanitize did not complete successfully may contain a vendor specific code that indicates a vendor specific reason)								
11:08	Status Field: Contains the value that was reported in the status code field for the completion queue entry, if any, for the Format NVM command associated with this event. If no completion queue entry was reported, then this field shall be cleared to 0h.								

5.14.1.TBD.1.9 Sanitize Start Event (Event Type 09h)

A Sanitize Start event shall be recorded in the Persistent Event Log at the start of a sanitize operation.

The Sanitize Start event shall set the Persistent Event Log Event Format Header:

- Event Type field to 09h and
- Event Type Revision field to 01h

Figure TBD + 16: Sanitize Start Event Data Format (Event Type 09h)

Bytes	Description
03:00	SANICAP: Contains the contents of the SANICAP field from the Identify Controller data structure.
07:04	Sanitize CDW10: Contains the value from command Dword 10 of the Sanitize command (refer to Figure 175).
11:08	Sanitize CDW11: Contains the value from command Dword 11 of the Sanitize command (refer to Figure 176).

5.14.1.TBD.1.10 Sanitize Completion Event (Event Type 0Ah)

A Sanitize Completion event shall be recorded in the Persistent Event Log at the completion of a sanitize operation.

The Sanitize Completion event shall set the Persistent Event Log Event Format Header:

- Event Type field to 0Ah and
- Event Type Revision field to 01h

Figure TBD + 17: Sanitize Completion Event Data Format (Event Type 0Ah)

Bytes	Description
1:0	Sanitize Progress: Contains the sanitize progress at the time of this event using the format specified for the SPROG field in the Sanitize Status log page (refer to section 5.14.1.9.2).
3:2	Sanitize Status: Contains the sanitize status for the time of this event using the format specified for the SSTAT field in the Sanitize Status log page. (e.g., the Global Data Erase bit indicates the status at the time of this event)
5:4	Completion Information: Contains a vendor specific value that may provide more information about the completion of the sanitize operation (e.g., if the sanitize did not complete successfully may contain a vendor specific code that indicates a vendor specific reason)
7:6	Reserved

5.14.1.TBD.1.11 Set Feature Event (Event Type 0Bh)

The Set Feature Event persists the data of a successful Set Features command. The event contains a Persistent Event Log Event header and the Set Feature Event data (refer to Figure TBD + 18).

The Set Feature Event shall set the Persistent Event Log Event Format Header:

- Event Type field to 0Bh; and
- Event Type Revision Field to 01h.

A Set Feature Event shall be recorded in the Persistent Event Log when the following criteria are met:

- A Set Feature command completes successfully;
- The Feature Identifier in that Set Feature command is supported to be logged in the Persistent Event Log; and
- There is a change to the controller settings for the Feature Identifier in that Set Feature command (i.e., the same setting is not set again).

A Set Feature Event may be recorded in the Persistent Event Log when there is no change to the controller settings for the Feature Identifier in that Set Feature command if the following criteria are met:

- A Set Feature command completes successfully; and
- The Feature Identifier in that Set Feature command is supported to be logged in the Persistent Event Log.

The Feature Identifiers that may be supported to be logged in the Persistent Event Log are shown in Figure 128 and Figure 129.

<Editor Note: When integrating with TP 4039, then the above Figure references of Figure 128 and Figure 129 would be Figure XA6, Figure XA7, Figure XB5, and Figure XB6>

The Command Dwords and Memory Buffer logged in the Set Feature Event data use the same formats as the formats defined by the Set Features and Get Features commands.

Figure TBD + 18: Set Feature Event Data Format

Bytes	Description										
03:00	<p>Set Feature Event Layout – Defines the number of Command Dwords and the amount of data in the Memory Buffer from the Set Features command associated with this event.</p> <table> <tr> <th>Bit</th><th>Definition</th></tr> <tr> <td>31:16</td><td>Memory Buffer Count: Defines the number of bytes from the memory buffer that are logged in the Memory Buffer field. A value of 0h indicates that the Memory Buffer field does not exist.</td></tr> <tr> <td>15:04</td><td>Reserved</td></tr> <tr> <td>03</td><td>Logged Command Completion Dword 0: If set to '1', then Dword 0 of the command completion for the Set Features command is included in the log. If cleared to '0', then Dword 0 of the command completion command for the Set Features command is not included in the log.</td></tr> <tr> <td>02:00</td><td>Dword Count: contains the number of consecutive Dwords starting with Dword 10 from the Set Feature command that are reported in the Command Dwords field. The values 0h and 7h are reserved.</td></tr> </table>	Bit	Definition	31:16	Memory Buffer Count: Defines the number of bytes from the memory buffer that are logged in the Memory Buffer field. A value of 0h indicates that the Memory Buffer field does not exist.	15:04	Reserved	03	Logged Command Completion Dword 0: If set to '1', then Dword 0 of the command completion for the Set Features command is included in the log. If cleared to '0', then Dword 0 of the command completion command for the Set Features command is not included in the log.	02:00	Dword Count: contains the number of consecutive Dwords starting with Dword 10 from the Set Feature command that are reported in the Command Dwords field. The values 0h and 7h are reserved.
Bit	Definition										
31:16	Memory Buffer Count: Defines the number of bytes from the memory buffer that are logged in the Memory Buffer field. A value of 0h indicates that the Memory Buffer field does not exist.										
15:04	Reserved										
03	Logged Command Completion Dword 0: If set to '1', then Dword 0 of the command completion for the Set Features command is included in the log. If cleared to '0', then Dword 0 of the command completion command for the Set Features command is not included in the log.										
02:00	Dword Count: contains the number of consecutive Dwords starting with Dword 10 from the Set Feature command that are reported in the Command Dwords field. The values 0h and 7h are reserved.										
(Dword Count * 4)+3: 4	Command Dwords: Contains a sequential list of Command Dwords from the Set Feature command starting with Command Dword 10. The number of entries in the list is specified by the Command Dword Count field. All non-reserved Command Dwords specified by the Set Feature command for the Feature Identifier shall be logged. The Command Dwords are ordered as defined by command format for the Admin Command Set and NVM Command in Figure 11.										
Data Buffer Count + (Dword Count * 4)+4: (Dword Count * 4)+4	<p>Memory Buffer: Contains the data in the memory buffer for the Set Features command.</p> <p>If the Memory Buffer Count field is cleared to a value on 0h, then this field does not exist in the logged event.</p>										
Data Buffer Count + (Dword Count * 4)+8: Data Buffer Count + (Dword Count * 4)+5	Command Completion Dword 0: If the Logged Command Completion Dword 0 bit is set to '1', then this field contains the Dword 0 value from the Set Features command completion. If the Logged Command Completion Dword 0 bit is cleared to '0', then this field is not logged.										

5.14.1.TBD.1.12 Telemetry Log Create Event (Event Type 0Ch)

A Telemetry Log Create event may be created if the controller determines that a host-initiated telemetry log (refer to section 5.14.1.7) or that a controller-initiated telemetry log (refer to section 5.14.1.8) which contains information about an NVM subsystem fault has been generated.

The Telemetry Log Create Event shall set the Persistent Event Log Event Format Header:

- Event Type field to 0Ch; and
- Event Type Revision Field to 01h.

Figure TBD + 19: Telemetry Log Create Event Data Format (Event Type 0Ch)

Bytes	Description
511:0	Telemetry Initiated Log: Contains a copy of the values from the first 512 bytes of the Host Initiated Log or the first 512 bytes of the Controller Initiated Log (refer to Figure 101 and Figure 102).

5.14.1.TBD.1.13 Thermal Excursion Event (Event Type 0Dh)

A Thermal Excursion event shall be recorded in the Persistent Event Log if the Composite Temperature has transitioned from a temperature that is less than:

- a) the WCTEMP, if any, (refer to Figure 111) to a temperature that is greater than or equal to the WCTEMP, if any; or
- b) the CCTEMP, if any, (refer to Figure 111), to a temperature that is greater than or equal to the CCTEMP, if any,

unless recording of the event causes a vendor specific frequency of threshold reports for this threshold to be exceeded.

A Thermal Excursion event may be recorded in the Persistent Event Log if the Composite Temperature has transitioned from a temperature:

- a) that is less than TMT1 (refer to section 5.21.1.16), if any, to a temperature that is greater than or equal to TMT1, if any (i.e., light throttling has started);
- b) that is less than TMT2 (refer to section 5.21.1.16), if any, to a temperature that is greater than or equal to TMT2, if any (i.e., heavy throttling has started);
- c) that is less than a vendor specific temperature where thermal throttling occurs due to self-throttling to a temperature that is greater than that vendor specific temperature (i.e., self-throttling has started);
- d) outside of a temperature threshold to a value that is within all temperature thresholds (i.e., the temperature returns to normal);
- e) at which thermal throttling is occurring to a temperature at which thermal throttling is stopped; or
- f) that is greater than an under temperature threshold (refer to section 5.21.1.4) to a temperature that is less than or equal to an under temperature threshold,

unless recording of the event causes a vendor specific frequency of threshold reports for this threshold to be exceeded.

The Thermal Excursion event shall set the Persistent Event Log Event Format Header;

- Event Type field to 0Dh; and
- Event Type Revision field to 01h.

Figure TBD + 20: Thermal Excursion Event Data Format (Event Type 0Dh)

Bytes	Field																						
04	Over Temperature: Contains the difference (delta) in degrees Kelvin between the temperature indicated in the threshold field and temperature measured at the time of the event.																						
021	Threshold: Contains an indicator for the temperature threshold crossing that is being reported. <table border="1"> <thead> <tr> <th>Value</th><th>Description</th></tr> </thead> <tbody> <tr> <td>1h</td><td>The Composite Temperature has transitioned from a temperature that is less than WCTEMP to a temperature that is greater than or equal to WCTEMP.</td></tr> <tr> <td>2h</td><td>The Composite Temperature has transitioned from a temperature that is less than CCTEMP to a temperature that is greater than or equal to CCTEMP.</td></tr> <tr> <td>3h</td><td>The Composite Temperature has transitioned from a temperature that is less than TMT1 to a temperature is greater than or equal to TMT1 (i.e., vendor specific thermal management actions that minimize the impact on performance, such as light throttling, have started).</td></tr> <tr> <td>4h</td><td>The Composite Temperature has transitioned from a temperature that is less than TMT2 to a temperature that is greater than or equal to TMT2 (i.e., vendor specific thermal management actions that may impact performance, such as heavy throttling, have started).</td></tr> <tr> <td>5h</td><td>The Composite Temperature has transitioned from a temperature where no vendor specific thermal management actions are taken to a temperature temperature that is greater than or equal to a vendor specific temperature at which vendor specific thermal management actions have started (e.g., self-throttling).</td></tr> <tr> <td>88h</td><td>The Composite Temperature has transitioned from a temperature that is greater than or equal to WCTEMP or is less than or equal to an under temperature threshold to a temperature that is between WCTEMP and that under temperature threshold (i.e., the temperature has transitioned to a normal temperature).</td></tr> <tr> <td>89h</td><td>The Composite Temperature has transitioned from a temperature that is greater than a temperature where thermal management actions are being performed and that is not greater than or equal to WCTEMP to a temperature where thermal management actions are stopped.</td></tr> <tr> <td>B0h</td><td>The Composite Temperature has transition from a temperature that is greater than an under temperature threshold to a temperature that is less than or equal to an under temperature threshold.</td></tr> <tr> <td>F0h to FFh</td><td>Vendor specific</td></tr> <tr> <td>All others</td><td>Reserved</td></tr> </tbody> </table>	Value	Description	1h	The Composite Temperature has transitioned from a temperature that is less than WCTEMP to a temperature that is greater than or equal to WCTEMP.	2h	The Composite Temperature has transitioned from a temperature that is less than CCTEMP to a temperature that is greater than or equal to CCTEMP.	3h	The Composite Temperature has transitioned from a temperature that is less than TMT1 to a temperature is greater than or equal to TMT1 (i.e., vendor specific thermal management actions that minimize the impact on performance, such as light throttling, have started).	4h	The Composite Temperature has transitioned from a temperature that is less than TMT2 to a temperature that is greater than or equal to TMT2 (i.e., vendor specific thermal management actions that may impact performance, such as heavy throttling, have started).	5h	The Composite Temperature has transitioned from a temperature where no vendor specific thermal management actions are taken to a temperature temperature that is greater than or equal to a vendor specific temperature at which vendor specific thermal management actions have started (e.g., self-throttling).	88h	The Composite Temperature has transitioned from a temperature that is greater than or equal to WCTEMP or is less than or equal to an under temperature threshold to a temperature that is between WCTEMP and that under temperature threshold (i.e., the temperature has transitioned to a normal temperature).	89h	The Composite Temperature has transitioned from a temperature that is greater than a temperature where thermal management actions are being performed and that is not greater than or equal to WCTEMP to a temperature where thermal management actions are stopped.	B0h	The Composite Temperature has transition from a temperature that is greater than an under temperature threshold to a temperature that is less than or equal to an under temperature threshold.	F0h to FFh	Vendor specific	All others	Reserved
Value	Description																						
1h	The Composite Temperature has transitioned from a temperature that is less than WCTEMP to a temperature that is greater than or equal to WCTEMP.																						
2h	The Composite Temperature has transitioned from a temperature that is less than CCTEMP to a temperature that is greater than or equal to CCTEMP.																						
3h	The Composite Temperature has transitioned from a temperature that is less than TMT1 to a temperature is greater than or equal to TMT1 (i.e., vendor specific thermal management actions that minimize the impact on performance, such as light throttling, have started).																						
4h	The Composite Temperature has transitioned from a temperature that is less than TMT2 to a temperature that is greater than or equal to TMT2 (i.e., vendor specific thermal management actions that may impact performance, such as heavy throttling, have started).																						
5h	The Composite Temperature has transitioned from a temperature where no vendor specific thermal management actions are taken to a temperature temperature that is greater than or equal to a vendor specific temperature at which vendor specific thermal management actions have started (e.g., self-throttling).																						
88h	The Composite Temperature has transitioned from a temperature that is greater than or equal to WCTEMP or is less than or equal to an under temperature threshold to a temperature that is between WCTEMP and that under temperature threshold (i.e., the temperature has transitioned to a normal temperature).																						
89h	The Composite Temperature has transitioned from a temperature that is greater than a temperature where thermal management actions are being performed and that is not greater than or equal to WCTEMP to a temperature where thermal management actions are stopped.																						
B0h	The Composite Temperature has transition from a temperature that is greater than an under temperature threshold to a temperature that is less than or equal to an under temperature threshold.																						
F0h to FFh	Vendor specific																						
All others	Reserved																						

5.14.1.TBD.1.14 Vendor Specific Event (Event Type DEh)

The content in this sub-clause is pulled from TP 4007, Persistent Event log.

The Vendor Specific Event (refer to Figure 91) contains a set of Vendor Specific Event Descriptors that describe an event that the vendor has determined is a significant event which should be reported to a host in the persistent event log and that is not described by any of the other persistent event log events.

The Vendor Specific Event Descriptors follow the format shown in Figure 92 and contain vendor specific data of the type indicated in the Vendor Specific Event Data Type field of the Vendor Specific Event Descriptor.

If a UUID Index is specified in the Get Log Page command (refer to section 5.14), then the controller shall return:

- a) Vendor specific events defined by the vendor identified by the specified UUID index; and
- b) Vendor specific events defined by the NVM subsystem manufacturer.

The controller shall set the Vendor Specific Event Format Header:

- a) Event Type field to DEh; and
- b) Event Type Revision field to 01h.

The Vendor Specific Event data is specified in Figure 91.

Figure 91: Vendor Specific Event Format (Event Type DEh)

Bytes	Description
M-1:0	Vendor Specific Event Descriptor 0: Contains the first vendor specific event descriptor (refer to Figure 92). Where M is the length of this vendor specific event descriptor.
...	
EL-VSIL-1: EL-VSIL-K	Vendor Specific Event Descriptor N: Contains the last vendor specific event descriptor (refer to Figure 92). Where K is the length of this vendor specific event descriptor.

The format of the Vendor Specific Event Descriptor is shown in Figure 92.

Figure 92: Vendor Specific Event Descriptor Format

Bytes	Description
01:00	Vendor Specific Event Code: Contains a vendor specific code that uniquely identifies the type of event that is described in the data that follows. All vendor specific events of the same type should report the same Vendor Specific Event Code field value.
02	Vendor Specific Event Data Type: Contains a code indicating the type of data reported in the Vendor Specific Event Data field (refer to Error! Reference source not found.).
03	UUID Index: UUID index (refer to <TP 4027> Figure TBD7) at the time of this event for the vendor that defined this event.
05:04	Vendor Specific Event Data Length: Contains the length in bytes of the Vendor Specific Event Data field.
M+5:06	Vendor Specific Event Data: Contains vendor specific data that is associated with this event and is of the type specified in the Vendor Specific Event Data Type field. Where M is the length of the vendor specific event data.

No further changes proposed in the Vendor Specific Event from TP 4007, following is all new content

5.14.1.TBD.1.15 TCG Defined Event (Event Type DFh)

The TCG Defined Event shall set the Persistent Event Log Event Format Header:

- Event Type field to DFh.

The Event Type Revision Field and the TCG Defined Event data are defined by TCG (refer to TCG Storage Interface Interactions specification).

No further changes in sub-clause 5.14.1

5.14.2 Feature Specific Information

Figure defines the Features that may be configured with Set Features and retrieved with Get Features. Figure 129 defines Features that are specific to the NVM Command Set. Some Features utilize a memory buffer to configure or return attributes for a Feature, whereas others only utilize a dword in the command or completion queue entry. Feature values that are not persistent across power cycles and resets are restored to their default values as part of a controller reset operation. The default value for each Feature is vendor specific and set by the manufacturer unless otherwise specified; it is not changeable. For more information on Features, including default, saveable, and current value definitions, refer to section 7.8.

There may be commands in execution when a Feature is changed. The new settings may or may not apply to commands already submitted for execution when the Feature is changed. Any commands submitted to a Submission Queue after a Set Features is successfully completed shall utilize the new settings for the associated Feature. To ensure that a Feature applies to all subsequent commands, commands being processed should be completed prior to issuing the Set Features command.

<Editor Note: When integrating with TP 4039, then the changes to Figure 128 and Figure 129 do not take effect as the information is moved by TP 4039 which is defined later in this document. >

Figure 128: Set Features – Feature Identifiers

Feature Identifier	O/M ⁶	Logged in Persistent Event log ⁶	Persistent Across Power Cycle and Reset ²	Uses Memory Buffer for Attributes	Description
00h					Reserved
01h	M	O	No	No	Arbitration
02h	M	NR	No	No	Power Management
03h	O	NR	Yes	Yes	LBA Range Type
04h	M	O	No	No	Temperature Threshold
05h	M	O	No	No	Error Recovery
06h	O	O	No	No	Volatile Write Cache
07h	M	O	No	No	Number of Queues
08h	NOTE 5	O	No	No	Interrupt Coalescing
09h	NOTE 5	O	No	No	Interrupt Vector Configuration
0Ah	M	O	No	No	Write Atomicity Normal
0Bh	M	NR	No	No	Asynchronous Event Configuration
0Ch	O	O	No	Yes	Autonomous Power State Transition
0Dh	O	O	No ³	No ⁴	Host Memory Buffer
0Eh	O	P	No	Yes	Timestamp
0Fh	O	O	No	No	Keep Alive Timer
10h	O	O	Yes	No	Host Controlled Thermal Management
11h	O	O	No	No	Non-Operational Power State Config
12h	O	O	Yes	No	Read Recovery Level Config
13h	O	O	Yes	Yes	Predictable Latency Mode Config
14h	O	P	Yes	No	Predictable Latency Mode Window
15h	O	O	No	No	LBA Status Information Report Interval Editorial Note: From TP 4016
16h	O	O	No	No	Host Behavior Support Editorial Note: From TP 4033

Figure 128: Set Features – Feature Identifiers

Feature Identifier	O/M ⁶	Logged in Persistent Event log ⁶	Persistent Across Power Cycle and Reset ²	Uses Memory Buffer for Attributes	Description
17h	O	O	Yes	No	Sanitize Config
18h	O	O	No	No	Endurance Group Event Configuration
19h to 77h					Reserved
78h to 7Fh		O	Refer to the NVMe Management Interface Specification for definition.		
80h to BFh					Command Set Specific (Reserved)
C0h to FFh		O			Vendor Specific ¹
NOTES: 1. The behavior of a controller in response to an inactive namespace ID to a vendor specific Feature Identifier is vendor specific. 2. This column is only valid if the feature is not saveable (refer to section 7.8). If the feature is saveable, then this column is not used and any feature may be configured to be saved across power cycles and reset. 3. The controller does not save settings for the Host Memory Buffer feature across power states and reset events, however, host software may restore the previous values. Refer to section 7.9. 4. The feature does not use a memory buffer for Set Features, but it does use a memory buffer for Get Features. Refer to section 7.9. 5. The feature is mandatory for NVMe over PCIe implementations. This feature is not supported for NVMe over Fabrics implementations. 6. O/M: O = Optional, M = Mandatory, P = Prohibited, NR = Not Recommended.					

Figure 129: Set Features, NVM Command Set Specific – Feature Identifiers

Feature Identifier	O/M ⁴	Logged in Persistent Event log ⁴	Persistent Across Power Cycle and Reset ¹	Uses Memory Buffer for Attributes	Description
80h	O	NR	Yes	No	Software Progress Marker
81h	O ²	O	No	Yes	Host Identifier
82h	O ³	O	No	No	Reservation Notification Mask
83h	O ³	O	Yes	No	Reservation Persistence
84h	O	O	No	No	Namespace Write Protection Config
85h to BFh					Reserved
NOTES: 1. This column is only valid if the feature is not saveable (refer to section 6.8). If the feature is saveable, then this column is not used and any feature may be configured to be saved across power cycles and reset. 2. Mandatory if reservations are supported as indicated in the Identify Controller data structure. 3. Mandatory if reservations are supported by the namespace as indicated by a non-zero value in the Reservation Capabilities (RESCAP) field in the Identify Namespace data structure. 4. O/M: O = Optional, M = Mandatory, P = Prohibited, NR = Not Recommended.					

<Editor Note> TP 4039 moved the support for Features to section 7.1. When this TP is integrated with TP 4039, then move the information added by this TP in figure 128 and figure 129 to match the following >

Modify a portion of section 7.1.1 (I/O Controller) from TP 4039 when this TP is integrated with TP 4039 as follows:

...

Figure XA6 and Figure XA7 define features that are mandatory, optional, and prohibited for an I/O controller.

Figure XA6: I/O Controller – Feature Support

Feature Name	Feature Support Requirements ¹	Logged in Persistent Event Log ¹
Arbitration	M	O
Power Management	M	NR
LBA Range Type	O	NR
Temperature Threshold	M	O
Error Recovery	M	O
Volatile Write Cache	O	O
Number of Queues	M	O
Interrupt Coalescing	NOTE 2	O
Interrupt Vector Configuration	NOTE 2	O
Write Atomicity Normal	M	O
Asynchronous Event Configuration	M	NR
Autonomous Power State Transition	O	O
Host Memory Buffer	O	O
Timestamp	O	P
Keep Alive Timer	O	O
Host Controlled Thermal Management	O	O
Non-Operational Power State Config	O	O
Read Recovery Level Config	O	O
Predictable Latency Mode Config	O	O
Predictable Latency Mode Window	O	P
LBA Status Information Attributes	O	O
Host Behavior Support	O	O
Sanitize Config	O	O
Endurance Group Event Configuration	O	O
Notes: 1. O = Optional, M = Mandatory, P = Prohibited, NR = Not Recommended 2. The feature is mandatory for NVMe over PCIe. This feature is not supported for NVMe over Fabrics.		

Figure XA7: I/O Controller – NVM Command Set Specific Feature Support

Feature Name	Feature Support Requirements ¹	Logged in Persistent Event Log ¹
Software Progress Marker	O	NR
Host Identifier	O ²	O
Reservation Notification Mask	O ³	O
Reservation Persistence	O ³	O
Namespace Write Protection Config	O	O
Notes: 1. O = Optional, M = Mandatory, P = Prohibited, NR = Not Recommended 2. Mandatory if reservations are supported as indicated in the Identify Controller data structure. Mandatory if reservations are supported by the namespace as indicated by a non-zero value in the Reservation Capabilities (RESCAP) field in the Identify Namespace data structure.		

Modify a portion of section 7.1.2 (I/O Controller) from TP 4039 when this TP is integrated with TP 4039 as follows:

Figure XB5: Administrative Controller – Feature Support

Feature Name	Feature Support Requirements 1	Logged in Persistent Event Log ¹
Arbitration	P	P
Power Management	O	NR
LBA Range Type	P	P
Temperature Threshold	O	O
Error Recovery	P	P
Volatile Write Cache	P	P
Number of Queues	P	P
Interrupt Coalescing	NOTE 2	NOTE 2
Interrupt Vector Configuration	NOTE 2	NOTE 2
Write Atomicity Normal	P	P
Asynchronous Event Configuration	O ³	NR
Autonomous Power State Transition	O	O
Host Memory Buffer	O	O
Timestamp	O	P
Keep Alive Timer	O	O
Host Controlled Thermal Management	O	O
Non-Operational Power State Config	O	O
Read Recovery Level Config	O	O
Predictable Latency Mode Config	O	O
Predictable Latency Mode Window	O	P
LBA Status Information Attributes	P	O
Host Behavior Support	O	O
Sanitize Config	O	O
Endurance Group Event Configuration	O	O
Notes: 1. O = Optional, M = Mandatory, P = Prohibited, NR = Not Recommended 2. The feature is optional for NVMe over PCIe. This feature is not supported for NVMe over Fabrics. Mandatory if Telemetry Log, Firmware Commit or SMART/Health Critical Warnings are supported.		

Figure XB6: Administrative Controller – NVM Command Set Specific Feature Support

Feature Name	Feature Support Requirements 1	Logged in Persistent Event Log ¹
Software Progress Marker	O	NR
Host Identifier	O ²	O
Reservation Notification Mask	O ³	O
Reservation Persistence	O ³	O
Namespace Write Protection Config	O	O
Notes: 1. O = Optional, M = Mandatory, P = Prohibited, NR = Not Recommended 2. Mandatory if reservations are supported as indicated in the Identify Controller data structure. Mandatory if reservations are supported by the namespace as indicated by a non-zero value in the Reservation Capabilities (RESCAP) field in the Identify Namespace data structure.		