



LEGAL NOTICE:

© **Copyright 2008 to 2022 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 2022 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
c/o VTM, Inc.
3855 SW 153rd Drive
Beaverton, OR 97003
USA
info@nvmexpress.org

Technical input submitted to NVM Express® is subject to the terms of the NVM Express® Participant’s agreement. Copyright © 2008 to 2022 NVM Express, Inc..

NVM Express® Technical Proposal (TP)

Technical Proposal ID	4142 Temperature Threshold Hysteresis
Revision Date	2022-12-13
Builds on Specification(s)	NVM Express Base Specification 2.0b
References	

Technical Proposal Author(s)

Name	Company
Andres Baez, Myron Loewen	Solidigm
Lee Prewitt	Microsoft

Technical Proposal Overview

This proposal intends to do the following:

- Adds a parameter to the Temperature Threshold feature that allows the host to specify hysteresis on the thermal event within the limits supported by the controller. The controller uses the parameter to determine the end of the thermal event.
- Adds a mechanism to notify the host once the thermal event has ended.

Revision History

Revision Date	Change Description
2022-03-23	Initial version
2022-03-24	Updated miscellaneous text according to the feedback provided. Modified text to clarify hysteresis usage when the Feature is supported for temperature sensors. Added text specifying the controller's behavior when the host uses hysteresis for a Threshold Temperature Select that is not supported for the Feature.
2022-03-31	Accepted all previous feedback changes. Addressed other language clarifications. Clarified TMPTHH field is the number of kelvins.
2022-04-06	Addressed feedback from Samsung
2022-04-19	Removed support bit as it is not needed to maintain backwards compatibility. Fixed wording to reflect this change. Consolidated the description of the functionality for the Critical Warning bits and SMART/Health asynchronous event as it relates to temperature threshold hysteresis in the Temperature Threshold Feature section.
2022-04-20	Reworded the description of the start and end of a temperature event. Added description to the SMART/Health Critical Warnings asynchronous event field to point to the Temperature Threshold Hysteresis section for its usage when the feature is supported.
2022-05-06	Removed the usage of the existing SMART/Health asynchronous event in favor of a new immediate event to notify the host of the end of a temperature threshold hysteresis event.
2022-05-17	Moved the log page requirement for immediate events into each individual event definition. Minor rewording in the Warning Composite Temperature Time field to add clarity
2022-06-16	Fixed typo in copied text that was converted into a table in the Identify data structure. Added field names along with the acronym for new field in various sections. Changed the Temperature Threshold Hysteresis Attributes field from 9 to 8 bits.
2022-06-23	Fixed typos missed from the last change.
2022-07-14	Addressed Samsung's feedback: <ul style="list-style-type: none">- Several editorial changes- Minor wording changes
2022-11-28	Integrated
2022-12-13	Editorial changes per Mike Allison and Andres Baez

Description of Specification Changes for NVM Express Specification 2.0b

New Features/Feature Enhancements/Required Changes:

- Adds a host-controlled temperature hysteresis parameter to the Temperature Threshold feature that the controller uses to determine the end of the thermal event.
- If the controller does not support the new parameter, or the host does not use the new parameter, the parameter has no effect on the Temperature Threshold feature.
- Adds the ability for the controller to generate an asynchronous event that indicates the end of a thermal event.
- References
 - Technical Proposal 4142

Markup Conventions:

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

Description of Specification Changes for NVM Express Base Specification 2.0b

5.2 Asynchronous Event Request command

Asynchronous events are used to notify host software of status, error, and health information as these events occur. To enable asynchronous events to be reported by the controller, host software needs to submit one or more Asynchronous Event Request commands to the controller. The controller specifies an event to the host by completing an Asynchronous Event Request command. Host software should expect that the controller may not execute the command immediately; the command should be completed when there is an event to be reported.

The Asynchronous Event Request command is submitted by host software to enable the reporting of asynchronous events from the controller. This command has no timeout. The controller posts a completion queue entry for this command when there is an asynchronous event to report to the host. If Asynchronous Event Request commands are outstanding when the controller is reset, then each of those commands is aborted and should not return a CQE.

All command specific fields are reserved.

Host software may submit multiple Asynchronous Event Request commands to reduce event reporting latency. The total number of simultaneously outstanding Asynchronous Event Request commands is limited by the Asynchronous Event Request Limit specified in the Identify Controller data structure in Figure 275.

Asynchronous events are grouped into event types. The event type is indicated in the Asynchronous Event Type field in Dword 0 of the completion queue entry for the Asynchronous Event Request command. When the controller posts a completion queue entry for an outstanding Asynchronous Event Request command and thus reports an asynchronous event, subsequent events of that event type are automatically masked by the controller until the host clears that event. Unless otherwise stated, an event is cleared by reading the log page associated with that event using the Get Log Page command (refer to section 5.16). If that log page is not accessible because media is not ready (i.e., the controller returns a status code of Admin Command Media Not Ready for the Get Log Page command requesting the log page), then the controller shall not post a completion queue entry for that asynchronous event until the controller is able to successfully return the log page that is required to be read to clear the asynchronous event.

The following event types are defined:

- a) **Error event:** Indicates a general error that is not associated with a specific command (refer to Figure 144). To clear this event, host software reads the Error Information log (refer to section 5.16.1.2) using the Get Log Page command with the Retain Asynchronous Event bit cleared to '0';
- b) **SMART / Health Status event:** Indicates a SMART or health status event (refer to Figure 145). To clear this event, host software reads the SMART / Health Information log (refer to section 5.16.1.3) using the Get Log Page command with the Retain Asynchronous Event bit cleared to '0'. The SMART / Health conditions that trigger asynchronous events may be configured in the Asynchronous Event Configuration feature using the Set Features command (refer to section 5.27.1.8);
- c) **Notice event:** Indicates a general event (refer to Figure 146). To clear this event, host software reads the appropriate log page as described in Figure 146. The conditions that trigger asynchronous events may be configured in the Asynchronous Event Configuration feature using the Set Features command (refer to section 5.27.1.8).
- d) **I/O Command Specific Status events:** Events that are specific to an I/O command (refer to Figure 147).
- e) **Immediate events:** Events that are only reported when an outstanding Asynchronous Event Request command exists at the time the event occurs (Refer to Figure 148). If the event occurs and there is no outstanding Asynchronous Event Request command, then the event shall not be reported; ~~and No log page is associated with these events. These events include:~~

~~A.—Normal NVM Subsystem Shutdown event;~~
and

- f) **Vendor Specific event:** Indicates a vendor specific event. To clear this event, host software reads the indicated vendor specific log page using the Get Log Page command with the Retain Asynchronous Event bit cleared to '0'.

The Sanitize Operation Completed With Unexpected Deallocation asynchronous event shall be supported if the controller supports the Sanitize Config feature (refer to section 5.27.1.19).

Asynchronous events are reported due to a new entry being added to a log page (e.g., Error Information log) or a status update (e.g., status in the SMART / Health log). A status change may be permanent (e.g., the media has become read only) or transient (e.g., the temperature reached or exceeded a threshold for a period of time). Host software should modify the event threshold or mask the event for transient and permanent status changes before issuing another Asynchronous Event Request command to avoid repeated reporting of asynchronous events.

If an event occurs for which reporting is enabled and there are no Asynchronous Event Request commands outstanding, the controller should retain the event information for that Asynchronous Event Type and use that information as a response to the next Asynchronous Event Request command that is received. If a Get Log Page command clears the event prior to receiving the Asynchronous Event Request command or if a power off condition occurs, then a notification is not sent. If multiple events of the same type occur that have identical responses to the Asynchronous Event Request command, then those events may be reported as a single response to an Asynchronous Event Request command. If multiple events occur that are of different types or have different responses to the Asynchronous Event Request command, then the controller should retain a queue of those events for reporting in responses to subsequent Asynchronous Event Request commands.

5.2.1 Command Completion

A completion queue entry is posted to the Admin Completion Queue if there is an asynchronous event to report to the host. Command specific status values associated with Asynchronous Event Request are defined in Figure 142.

Figure 145: Asynchronous Event Information – SMART / Health Status

Value	Description
00h	NVM subsystem Reliability: NVM subsystem reliability has been compromised. This may be due to significant media errors, an internal error, the media being placed in read only mode, or a volatile memory backup device failing. This status value shall not be used if the read-only condition on the media is due to a change in the write protection state of a namespace (refer to section 8.12.1).
01h	Temperature Threshold: A temperature is greater than or equal to an over temperature threshold or less than or equal to an under temperature threshold (refer to section 5.27.1.3).
02h	Spare Below Threshold: Available spare capacity has fallen below the threshold.
03h to FFh	Reserved

Figure 1: Asynchronous Event Information – Immediate

Value	Description
00h	NVM Subsystem Normal Shutdown: This controller has started performing a normal NVM Subsystem Shutdown that is due to: <ul style="list-style-type: none"> the value 4E726D6Ch ("Nrml") has been written to an NSS.NCCR register within the NVM subsystem or Domain; or an NVMe-MI Shutdown command (refer to the NVMe Express Management Interface Specification) being processed. Refer to section Error! Reference source not found..
01h	Temperature Threshold Hysteresis Recovery: Indicates the end of a temperature threshold hysteresis event (Refer to section 5.27.1.3.TBD).
04h 02h to FFh	Reserved

5.16.1.3 SMART / Health Information (Log Identifier 02h)

This log page is used to provide SMART and general health information. The information provided is over the life of the controller and is retained across power cycles. To request the controller log page, the namespace identifier specified is FFFFFFFFh or 0h. For compatibility with implementations compliant with NVM Express Base Specification revision 1.4 and earlier, hosts should use a namespace identifier of FFFFFFFFh to request the controller log page. The controller may also support requesting the log page on a per namespace basis, as indicated by bit 0 of the LPA field in the Identify Controller data structure in Figure 275.

If the log page is not supported on a per namespace basis, specifying a namespace identifier other than 0h or FFFFFFFFh should abort the command with a status code of Invalid Field in Command. If the controller does not abort the command, then the controller returns the controller log page. There is no namespace specific information defined in the SMART / Health Information log page in this revision of the specification, thus the controller log page and namespaces specific log page contain identical information.

Critical warnings regarding the health of the NVM subsystem may be indicated via an asynchronous event notification to the host. The warnings that results in an asynchronous event notification to the host are configured using the Set Features command; refer to section 5.27.1.8.

Performance may be calculated using parameters returned as part of the SMART / Health Information log. Specifically, the number of Read or Write commands, the amount of data read or written, and the amount of controller busy time enables both I/Os per second and bandwidth to be calculated.

The log page returned is defined in Figure 207.

Figure 207: Get Log Page – SMART / Health Information Log

Bytes	Description																
00	Critical Warning: This field indicates critical warnings for the state of the controller. Each bit corresponds to a critical warning type; multiple bits may be set to '1'. If a bit is cleared to '0', then that critical warning does not apply. Critical warnings may result in an asynchronous event notification to the host. Bits in this field represent the current associated state and are not persistent.																
	<table><tr><th>Bits</th><th>Definition</th></tr><tr><td>7:6</td><td>Reserved</td></tr><tr><td>5</td><td>If set to '1', then the Persistent Memory Region has become read-only or unreliable (refer to section 8.14).</td></tr><tr><td>4</td><td>If set to '1', then the volatile memory backup device has failed. This field is only valid if the controller has a volatile memory backup solution.</td></tr><tr><td>3</td><td>If set to '1', then all of the media has been placed in read only mode. The controller shall not set this bit to '1' if the read-only condition on the media is a result of a change in the write protection state of a namespace (refer to section 8.12.1).</td></tr><tr><td>2</td><td>If set to '1', then the NVM subsystem reliability has been degraded due to significant media related errors or any internal error that degrades NVM subsystem reliability.</td></tr><tr><td>1</td><td>If set to '1', then a temperature is: <div>a) greater than or equal to an over temperature threshold; or b) less than or equal to an under temperature threshold, (refer to section 5.27.1.3).</div></td></tr><tr><td>0</td><td>If set to '1', then the available spare capacity has fallen below the threshold.</td></tr></table>	Bits	Definition	7:6	Reserved	5	If set to '1', then the Persistent Memory Region has become read-only or unreliable (refer to section 8.14).	4	If set to '1', then the volatile memory backup device has failed. This field is only valid if the controller has a volatile memory backup solution.	3	If set to '1', then all of the media has been placed in read only mode. The controller shall not set this bit to '1' if the read-only condition on the media is a result of a change in the write protection state of a namespace (refer to section 8.12.1).	2	If set to '1', then the NVM subsystem reliability has been degraded due to significant media related errors or any internal error that degrades NVM subsystem reliability.	1	If set to '1', then a temperature is: <div>a) greater than or equal to an over temperature threshold; or b) less than or equal to an under temperature threshold, (refer to section 5.27.1.3).</div>	0	If set to '1', then the available spare capacity has fallen below the threshold.
	Bits	Definition															
	7:6	Reserved															
	5	If set to '1', then the Persistent Memory Region has become read-only or unreliable (refer to section 8.14).															
	4	If set to '1', then the volatile memory backup device has failed. This field is only valid if the controller has a volatile memory backup solution.															
	3	If set to '1', then all of the media has been placed in read only mode. The controller shall not set this bit to '1' if the read-only condition on the media is a result of a change in the write protection state of a namespace (refer to section 8.12.1).															
	2	If set to '1', then the NVM subsystem reliability has been degraded due to significant media related errors or any internal error that degrades NVM subsystem reliability.															
	1	If set to '1', then a temperature is: <div>a) greater than or equal to an over temperature threshold; or b) less than or equal to an under temperature threshold, (refer to section 5.27.1.3).</div>															
0	If set to '1', then the available spare capacity has fallen below the threshold.																
02:01	Composite Temperature: Contains 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 manner in which this value is computed is implementation specific and may not represent the actual temperature of any physical point in the NVM subsystem. The value of this field may be used to trigger an asynchronous event (refer to section 5.27.1.3).																
	Warning and critical overheating composite temperature threshold values are reported by the WCTEMP and CCTEMP fields in the Identify Controller data structure in Figure 275.																

Figure 207: Get Log Page – SMART / Health Information Log

Bytes	Description
191:176	Number of Error Information Log Entries: Contains the number of Error Information log entries over the life of the controller.
195:192	<p>Warning Composite Temperature Time: If the Temperature Threshold Hysteresis Attributes (TMPTHHA) field is cleared to 0h (refer to Figure 275), then this field contains the amount of time in minutes that the controller is operational and the Composite Temperature is greater than or equal to the Warning Composite Temperature Threshold (WCTEMP) field and less than the Critical Composite Temperature Threshold (CCTEMP) field in the Identify Controller data structure in Figure 275.</p> <p>If the Temperature Threshold Maximum Hysteresis (TMPTHMH) and Temperature Threshold Hysteresis (TMPTHH) fields are non-zero, then hysteresis time is included in the time specified in this field (refer to section 5.27.1.3.TBD).</p> <p>If the value of the WCTEMP or CCTEMP field is 0h, then this field is always cleared to 0h regardless of the Composite Temperature value.</p>
199:196	<p>Critical Composite Temperature Time: Contains the amount of time in minutes that the controller is operational and the Composite Temperature is greater than or equal to the Critical Composite Temperature Threshold (CCTEMP) field in the Identify Controller data structure in Figure 275.</p> <p>If the value of the CCTEMP field is 0h, then this field is always cleared to 0h regardless of the Composite Temperature value.</p>
201:200	Temperature Sensor 1: Contains the current temperature reported by temperature sensor 1. This field is defined by Figure 208.
203:202	Temperature Sensor 2: Contains the current temperature reported by temperature sensor 2. This field is defined by Figure 208.
205:204	Temperature Sensor 3: Contains the current temperature reported by temperature sensor 3. This field is defined by Figure 208.
207:206	Temperature Sensor 4: Contains the current temperature reported by temperature sensor 4. This field is defined by Figure 208.
209:208	Temperature Sensor 5: Contains the current temperature reported by temperature sensor 5. This field is defined by Figure 208.
211:210	Temperature Sensor 6: Contains the current temperature reported by temperature sensor 6. This field is defined by Figure 208.
213:212	Temperature Sensor 7: Contains the current temperature reported by temperature sensor 7. This field is defined by Figure 208.
215:214	Temperature Sensor 8: Contains the current temperature reported by temperature sensor 8. This field is defined by Figure 208.
219:216	Thermal Management Temperature 1 Transition Count: Contains the number of times the controller transitioned to lower power active power states or performed vendor specific thermal management actions while minimizing the impact on performance in order to attempt to reduce the Composite Temperature because of the host controlled thermal management feature (refer to section 8.15.5) (i.e., the Composite Temperature rose above the Thermal Management Temperature 1). This counter shall not wrap once the value FFFFFFFFh is reached. A value of 0h, indicates that this transition has never occurred or this field is not implemented.
223:220	Thermal Management Temperature 2 Transition Count: Contains the number of times the controller transitioned to lower power active power states or performed vendor specific thermal management actions regardless of the impact on performance (e.g., heavy throttling) in order to attempt to reduce the Composite Temperature because of the host controlled thermal management feature (refer to section 8.15.5) (i.e., the Composite Temperature rose above the Thermal Management Temperature 2). This counter shall not wrap once the value FFFFFFFFh is reached. A value of 0h, indicates that this transition has never occurred or this field is not implemented.

Figure 207: Get Log Page – SMART / Health Information Log

Bytes	Description
227:224	Total Time For Thermal Management Temperature 1: Contains the number of seconds that the controller had transitioned to lower power active power states or performed vendor specific thermal management actions while minimizing the impact on performance in order to attempt to reduce the Composite Temperature because of the host controlled thermal management feature (refer to section 8.15.5). This counter shall not wrap once the value FFFFFFFFh is reached. A value of 0h, indicates that this transition has never occurred or this field is not implemented.
231:228	Total Time For Thermal Management Temperature 2: Contains the number of seconds that the controller had transitioned to lower power active power states or performed vendor specific thermal management actions regardless of the impact on performance (e.g., heavy throttling) in order to attempt to reduce the Composite Temperature because of the host controlled thermal management feature (refer to section 8.15.5). This counter shall not wrap once the value FFFFFFFFh is reached. A value of 0h, indicates that this transition has never occurred or this field is not implemented.
511:232	Reserved

5.17.2 Identify Data Structures

5.17.2.1 Identify Controller data structure (CNS 01h)

The Identify Controller data structure (refer to Figure 275) is returned to the host for the controller processing the command.

Figure 275: Identify – Identify Controller Data Structure, I/O Command Set Independent

95:92	M	M	M	<p>Optional Asynchronous Events Supported (OAES): This field indicates the optional asynchronous events supported by the controller. A controller shall not send optional asynchronous events before they are enabled by host software.</p> <p>Bit 31 is set to '1' if the controller supports sending Discovery Log Page Change Notifications. If cleared to '0', then the controller does not support the Discovery Log Page Change Notification events.</p> <p>Bits 30:28 are reserved.</p> <p>Bit 27 is set to '1' if the controller supports the Zone Descriptor Changed Notices event and the associated Changed Zone List log page (refer to the Zoned Namespace Command Set specification). If cleared to '0', then the controller does not support the Zone Descriptor Changed Notices event nor the associated Changed Zone List log page.</p> <p>Bits 26:16 are reserved.</p> <p>Bit 15 is set to '1' if the controller supports the Normal NVM Subsystem Shutdown event. If cleared to '0', then the controller does not support the Normal NVM Subsystem Shutdown event.</p> <p>Bit 14 is set to '1' if the controller supports the Endurance Group Event Aggregate Log Page Change Notices event. If cleared to '0', then the controller does not support the Endurance Group Event Aggregate Log Page Change Notices event.</p> <p>Bit 13 is set to '1' if the controller supports the LBA Status Information Alert Notices event (refer to the NVM Command Set specification). If cleared to '0', then the controller does not support the LBA Status Information Alert Notices event.</p> <p>Bit 12 is set to '1' if the controller supports the Predictable Latency Event Aggregate Log Change Notices event. If cleared to '0', then the controller does not support the Predictable Latency Event Aggregate Log Change Notices event.</p>
-------	---	---	---	--

Figure 275: Identify – Identify Controller Data Structure, I/O Command Set Independent

			<p>Bit 11 is set to '1' if the controller supports sending Asymmetric Namespace Access Change Notices. If cleared to '0', then the controller does not support the Asymmetric Namespace Access Change Notices event.</p> <p>Bit 10 is reserved.</p> <p>Bit 9 is set to '1' if the controller supports the Firmware Activation Notices event. If cleared to '0', then the controller does not support the Firmware Activation Notices event.</p> <p>Bit 8 is set to '1' if the controller supports the Namespace Attribute Notices event and the associated Changed Namespace List log page. If cleared to '0', then the controller does not support the Namespace Attribute Notices event nor the associated Changed Namespace List log page.</p> <p>Bits 7:0 are reserved.</p>																								
			<table><tr><th>Bits</th><th>Definition</th></tr><tr><td>31</td><td>Discovery Log Page Change Notification: If set to '1', then the controller supports sending Discovery Log Page Change Notifications. If cleared to '0', then the controller does not support the Discovery Log Page Change Notification events.</td></tr><tr><td>30:28</td><td>Reserved</td></tr><tr><td>27</td><td>Zone Descriptor Changed Notices: If set to '1', then the controller supports the Zone Descriptor Changed Notices event and the associated Changed Zone List log page (refer to the Zoned Namespace Command Set specification). If cleared to '0', then the controller does not support the Zone Descriptor Changed Notices event nor the associated Changed Zone List log page.</td></tr><tr><td>26:17</td><td>Reserved</td></tr><tr><td>16</td><td>Temperature Threshold Hysteresis Recovery: If set to '1', then the controller supports the Temperature Threshold Hysteresis Recovery event. If cleared to '0', then the controller does not support the Temperature Threshold Hysteresis Recovery event. No log page is associated with this event.</td></tr><tr><td>15</td><td>Normal NVM Subsystem Shutdown: If set to '1', then the controller supports the Normal NVM Subsystem Shutdown event. If cleared to '0', then the controller does not support the Normal NVM Subsystem Shutdown event. <Sentence moved from section 5.2> No log page is associated with this event.</td></tr><tr><td>14</td><td>Endurance Group Event Aggregate Log Page Change Notices: If set to '1', then the controller supports the Endurance Group Event Aggregate Log Page Change Notices event. If cleared to '0', then the controller does not support the Endurance Group Event Aggregate Log Page Change Notices event.</td></tr><tr><td>13</td><td>LBA Status Information Alert Notices: If set to '1', then the controller supports the LBA Status Information Alert Notices event (refer to the NVM Command Set specification). If cleared to '0', then the controller does not support the LBA Status Information Alert Notices event.</td></tr><tr><td>12</td><td>Predictable Latency Event Aggregate Log Change Notices: If set to '1', then controller supports the Predictable Latency Event Aggregate Log Change Notices event. If cleared to '0', then the controller does not support the Predictable Latency Event Aggregate Log Change Notices event.</td></tr><tr><td>11</td><td>Asymmetric Namespace Access Change Notices: If set to '1', then the controller supports sending Asymmetric Namespace Access Change Notices. If cleared to '0', then the controller does not support the Asymmetric Namespace Access Change Notices event.</td></tr><tr><td>10</td><td>Reserved</td></tr></table>	Bits	Definition	31	Discovery Log Page Change Notification: If set to '1', then the controller supports sending Discovery Log Page Change Notifications. If cleared to '0', then the controller does not support the Discovery Log Page Change Notification events.	30:28	Reserved	27	Zone Descriptor Changed Notices: If set to '1', then the controller supports the Zone Descriptor Changed Notices event and the associated Changed Zone List log page (refer to the Zoned Namespace Command Set specification). If cleared to '0', then the controller does not support the Zone Descriptor Changed Notices event nor the associated Changed Zone List log page.	26:17	Reserved	16	Temperature Threshold Hysteresis Recovery: If set to '1', then the controller supports the Temperature Threshold Hysteresis Recovery event. If cleared to '0', then the controller does not support the Temperature Threshold Hysteresis Recovery event. No log page is associated with this event.	15	Normal NVM Subsystem Shutdown: If set to '1', then the controller supports the Normal NVM Subsystem Shutdown event. If cleared to '0', then the controller does not support the Normal NVM Subsystem Shutdown event. <Sentence moved from section 5.2> No log page is associated with this event.	14	Endurance Group Event Aggregate Log Page Change Notices: If set to '1', then the controller supports the Endurance Group Event Aggregate Log Page Change Notices event. If cleared to '0', then the controller does not support the Endurance Group Event Aggregate Log Page Change Notices event.	13	LBA Status Information Alert Notices: If set to '1', then the controller supports the LBA Status Information Alert Notices event (refer to the NVM Command Set specification). If cleared to '0', then the controller does not support the LBA Status Information Alert Notices event.	12	Predictable Latency Event Aggregate Log Change Notices: If set to '1', then controller supports the Predictable Latency Event Aggregate Log Change Notices event. If cleared to '0', then the controller does not support the Predictable Latency Event Aggregate Log Change Notices event.	11	Asymmetric Namespace Access Change Notices: If set to '1', then the controller supports sending Asymmetric Namespace Access Change Notices. If cleared to '0', then the controller does not support the Asymmetric Namespace Access Change Notices event.	10	Reserved
Bits	Definition																										
31	Discovery Log Page Change Notification: If set to '1', then the controller supports sending Discovery Log Page Change Notifications. If cleared to '0', then the controller does not support the Discovery Log Page Change Notification events.																										
30:28	Reserved																										
27	Zone Descriptor Changed Notices: If set to '1', then the controller supports the Zone Descriptor Changed Notices event and the associated Changed Zone List log page (refer to the Zoned Namespace Command Set specification). If cleared to '0', then the controller does not support the Zone Descriptor Changed Notices event nor the associated Changed Zone List log page.																										
26:17	Reserved																										
16	Temperature Threshold Hysteresis Recovery: If set to '1', then the controller supports the Temperature Threshold Hysteresis Recovery event. If cleared to '0', then the controller does not support the Temperature Threshold Hysteresis Recovery event. No log page is associated with this event.																										
15	Normal NVM Subsystem Shutdown: If set to '1', then the controller supports the Normal NVM Subsystem Shutdown event. If cleared to '0', then the controller does not support the Normal NVM Subsystem Shutdown event. <Sentence moved from section 5.2> No log page is associated with this event.																										
14	Endurance Group Event Aggregate Log Page Change Notices: If set to '1', then the controller supports the Endurance Group Event Aggregate Log Page Change Notices event. If cleared to '0', then the controller does not support the Endurance Group Event Aggregate Log Page Change Notices event.																										
13	LBA Status Information Alert Notices: If set to '1', then the controller supports the LBA Status Information Alert Notices event (refer to the NVM Command Set specification). If cleared to '0', then the controller does not support the LBA Status Information Alert Notices event.																										
12	Predictable Latency Event Aggregate Log Change Notices: If set to '1', then controller supports the Predictable Latency Event Aggregate Log Change Notices event. If cleared to '0', then the controller does not support the Predictable Latency Event Aggregate Log Change Notices event.																										
11	Asymmetric Namespace Access Change Notices: If set to '1', then the controller supports sending Asymmetric Namespace Access Change Notices. If cleared to '0', then the controller does not support the Asymmetric Namespace Access Change Notices event.																										
10	Reserved																										

Figure 275: Identify – Identify Controller Data Structure, I/O Command Set Independent

				9	Firmware Activation Notices: If set to '1', then the controller supports the Firmware Activation Notices event. If cleared to '0', then the controller does not support the Firmware Activation Notices event.
				8	Namespace Attribute Notices: If set to '1', then the controller supports the Namespace Attribute Notices event and the associated Changed Namespace List log page. If cleared to '0', then the controller does not support the Namespace Attribute Notices event nor the associated Changed Namespace List log page.
				7:0	Reserved
...	
Bytes	I/O ¹	Admin ¹	Disc ¹	Description	
Controller Capabilities and Features					
367:358				Reserved	
383:368	O	R	R	Max Endurance Group Capacity (MEGCAP): This field indicates the maximum capacity of a single Endurance Group. If this field is cleared to 0h, the NVM subsystem does not report a maximum Endurance Group Capacity value.	
384	O	O	O	Temperature Threshold Hysteresis Attributes (TMPTHHA): This field indicates attributes related to temperature threshold hysteresis. Refer to section 5.27.1.3.TBD.	
				Bits	Description
				7:3	Reserved
				2:0	Temperature Threshold Maximum Hysteresis (TMPTHMH): This field indicates the maximum temperature hysteresis value in Kelvins that is supported by the controller. This is the absolute value of the difference. If the controller does not support this attribute, this field shall be cleared to 000b.
511:384385				Reserved	
NVM Command Set Attributes					
512	M	M	R	Submission Queue Entry Size (SQES): This field defines the required and maximum I/O Submission Queue entry size. Bits 7:4 define the maximum I/O Submission Queue entry size when using the NVM Command Set. This value is greater than or equal to the required SQ entry size (i.e., bits 3:0 in this field). The value is in bytes and is reported as a power of two (2^n). The recommended value is 6, corresponding to a standard SQ entry size of 64 bytes. Controllers that implement proprietary extensions may support a larger value. Bits 3:0 define the required (i.e., minimum) I/O Submission Queue entry size. This is the minimum entry size that may be used. The value is in bytes and is reported as a power of two (2^n). The required value shall be 6, corresponding to 64.	

5.27.1.3 Temperature Threshold (Feature Identifier 04h)

A controller may report up to nine temperature values in the SMART / Health Information log (i.e., the Composite Temperature and Temperature Sensor 1 through Temperature Sensor 8; refer to Figure 207). Associated with each implemented temperature sensor is an over temperature threshold and an under temperature threshold. When a temperature is greater than or equal to its corresponding over temperature threshold or less than or equal to its corresponding under temperature threshold, then bit one of the Critical Warning field in the SMART / Health Information Log (refer to section 5.16.1.3) is set to **one '1'**. This may trigger an asynchronous event.

The over temperature threshold feature shall be implemented for Composite Temperature. The under temperature threshold Feature shall be implemented for Composite Temperature if a non-zero Warning

Composite Temperature Threshold (WCTEMP) field value is reported in the Identify Controller data structure (refer to Figure 275). The over temperature threshold and under temperature threshold features shall be implemented for all implemented temperature sensors (i.e., all Temperature Sensor fields that report a non-zero value).

The default value of the over temperature threshold feature for Composite Temperature is the value in the Warning Composite Temperature Threshold (WCTEMP) field in the Identify Controller data structure if WCTEMP is non-zero; otherwise, the default value is implementation specific. The default value of the under temperature threshold feature for Composite Temperature is implementation specific. The default value of the over temperature threshold for all implemented temperature sensors is FFFFh. The default value of the under temperature threshold for all implemented temperature sensors is 0h.

If a Get Features command is submitted for this feature, the temperature threshold selected by Command Dword 11 is returned in Dword 0 of the completion queue entry for that command.

5.27.1.3.TBD Temperature Threshold Hysteresis

This feature allows the host to specify the temperature hysteresis that a controller shall use to determine the end of a temperature threshold event.

If the controller supports the use of the Temperature Threshold Hysteresis feature as indicated by a non-zero value of the TMPTHMH field (refer to Figure 275), then the controller shall:

- support this feature for the Composite Temperature value (refer to Figure 320); and
- support the Temperature Threshold Hysteresis Recovery event (refer to Figure 275).

If hysteresis is supported for Threshold Temperature Select values other than Composite Temperature (i.e., 0h) (refer to Figure 320); then, associated with each implemented temperature sensor is an over-temperature threshold and an under-temperature threshold with a hysteresis value. If the host specifies hysteresis for a Threshold Temperature Select that is not supported, the controller shall abort the command with a status of Invalid Field in Command.

An over-temperature threshold hysteresis event begins when the value of a temperature field (refer to the temperature values reported as described in section 5.27.1.3) transitions from less than its corresponding temperature threshold to greater than or equal to that temperature threshold. That over-temperature threshold hysteresis event ends when the value of that temperature field transitions from greater than or equal to that temperature threshold to less than that temperature threshold minus the value of the Temperature Threshold Hysteresis (TMPTHH) field.

An under-temperature threshold hysteresis event begins when the value of a temperature field transitions from greater than its corresponding temperature threshold to less than or equal to that temperature threshold. That under-temperature threshold hysteresis event ends when the value of that temperature field transitions from less than or equal to that temperature threshold to greater than that temperature threshold plus the value of the Temperature Threshold Hysteresis (TMPTHH) field.

When a temperature threshold hysteresis event has started, bit 1 of the Critical Warning field in the SMART / Health Information log page (refer to section 5.16.1.3) is set to '1'. This may trigger an initial asynchronous event.

At the end of a temperature threshold hysteresis event, a Temperature Threshold Hysteresis Recovery event shall be triggered (refer to Figure 326), bit 1 of the Critical Warning field in the SMART / Health Information log page (refer to section 5.16.1.3) shall be cleared to '0' and the Warning Composite Temperature Time field shall stop accumulating the number of minutes.

If the value of the Temperature Threshold Hysteresis field is larger than the value supported by the controller as specified by the TMPTHMH field in Figure 275, the controller shall abort the command with a status code of Invalid Field in Command.

Figure 320: Temperature Threshold – Command Dword 11

Bits	Description																								
31:225	Reserved																								
24:22	Temperature Threshold Hysteresis (TMPTHH): This field indicates the temperature hysteresis (i.e., the number of kelvins the controller uses to determine the end of the over temperature or under temperature event). If this field is cleared to 000b, hysteresis as defined for this Feature does not apply. If the Temperature Threshold Maximum Hysteresis (TMPTHMH) field in Figure 275 is cleared to 0h, this field shall be cleared to 000b.																								
21:20	Threshold Type Select (THSEL): This field selects the threshold type that is modified by a Set Features command and whose threshold value is returned by a Get Features command. <table border="1"> <thead> <tr> <th>Value</th><th>Description</th></tr> </thead> <tbody> <tr> <td>00b</td><td>Over Temperature Threshold</td></tr> <tr> <td>01b</td><td>Under Temperature Threshold</td></tr> <tr> <td>10b to 11b</td><td>Reserved</td></tr> </tbody> </table>	Value	Description	00b	Over Temperature Threshold	01b	Under Temperature Threshold	10b to 11b	Reserved																
Value	Description																								
00b	Over Temperature Threshold																								
01b	Under Temperature Threshold																								
10b to 11b	Reserved																								
19:16	Threshold Temperature Select (TMPSEL): This field selects the temperature whose threshold is modified by a Set Features command and whose threshold value is returned by a Get Features command. <table border="1"> <thead> <tr> <th>Value</th><th>Description</th></tr> </thead> <tbody> <tr> <td>0h</td><td>Composite Temperature</td></tr> <tr> <td>1h</td><td>Temperature Sensor 1</td></tr> <tr> <td>2h</td><td>Temperature Sensor 2</td></tr> <tr> <td>3h</td><td>Temperature Sensor 3</td></tr> <tr> <td>4h</td><td>Temperature Sensor 4</td></tr> <tr> <td>5h</td><td>Temperature Sensor 5</td></tr> <tr> <td>6h</td><td>Temperature Sensor 6</td></tr> <tr> <td>7h</td><td>Temperature Sensor 7</td></tr> <tr> <td>8h</td><td>Temperature Sensor 8</td></tr> <tr> <td>9h to Eh</td><td>Reserved</td></tr> <tr> <td>Fh</td><td>All implemented temperature sensors in a Set Features command. Reserved in a Get Features command.</td></tr> </tbody> </table>	Value	Description	0h	Composite Temperature	1h	Temperature Sensor 1	2h	Temperature Sensor 2	3h	Temperature Sensor 3	4h	Temperature Sensor 4	5h	Temperature Sensor 5	6h	Temperature Sensor 6	7h	Temperature Sensor 7	8h	Temperature Sensor 8	9h to Eh	Reserved	Fh	All implemented temperature sensors in a Set Features command. Reserved in a Get Features command.
Value	Description																								
0h	Composite Temperature																								
1h	Temperature Sensor 1																								
2h	Temperature Sensor 2																								
3h	Temperature Sensor 3																								
4h	Temperature Sensor 4																								
5h	Temperature Sensor 5																								
6h	Temperature Sensor 6																								
7h	Temperature Sensor 7																								
8h	Temperature Sensor 8																								
9h to Eh	Reserved																								
Fh	All implemented temperature sensors in a Set Features command. Reserved in a Get Features command.																								
15:00	Temperature Threshold (TMPTH): Indicates the threshold value for the temperature sensor and threshold type specified in Kelvins.																								

5.27.1.8 Asynchronous Event Configuration (Feature Identifier 0Bh)

This Feature controls the events that trigger an asynchronous event notification to the host. This Feature may be used to disable reporting events in the case of a persistent condition (refer to section 5.2). If the condition for an event is true when the corresponding notice is enabled, then an event is sent to the host. The attributes are specified in Command Dword 11.

If a Get Features command is submitted for this Feature, the attributes specified in Figure 326 are returned in Dword 0 of the completion queue entry for that command.

Figure 326: Asynchronous Event Configuration – Command Dword 11

Bits	Description
31	Discovery Log Page Change Notification: This bit indicates that the Discovery controller reports Discovery Log Page Change Notifications. If set to '1', the Discovery controller shall send a notification if Discovery Log Page changes occur.
30:28	Reserved
27	Zone Descriptor Changed Notices²: I/O Command Set specific definition.
26:17	Reserved

Figure 326: Asynchronous Event Configuration – Command Dword 11

Bits	Description
16	Temperature Threshold Hysteresis Recovery: This bit determines whether an asynchronous event notification is sent to the host at the end of a temperature threshold hysteresis event (refer to section 5.27.1.3.TBD). If this bit is set to '1', then the Temperature Threshold Hysteresis Recovery event is sent to the host if an outstanding Asynchronous Event Request command exists at the time this condition occurs. If this bit is cleared to '0', then the controller shall not send the Temperature Threshold Hysteresis Recovery event to the host.
15	Normal NVM Subsystem Shutdown: This bit determines whether an asynchronous event notification is sent to the host when the NVM subsystem has started performing a normal shutdown due to an NVM Subsystem Shutdown (refer to Figure 148). If this bit is set to '1', then the Normal NVM Subsystem Shutdown event is sent to the host if an outstanding Asynchronous Event Request command exists at the time this condition occurs. If this bit is cleared to '0', then the controller shall not send the Normal NVM Subsystem Shutdown event to the host.
14	Endurance Group Event Aggregate Log Change Notices: This bit determines whether an asynchronous event notification is sent to the host when an event entry for an Endurance Group (refer to section 3.2.3) has been added to the Endurance Group Event Aggregate log (refer to section 5.16.1.15). If this bit is set to '1', then the Endurance Group Event Aggregate Log Change event is sent to the host when this condition occurs. If this bit is cleared to '0', then the controller shall not send the Endurance Group Event Aggregate Log Change event to the host. If Endurance Groups are not supported and this bit is set to '1', then the Set Features command shall be aborted with a status of Invalid Field in Command.
13	LBA Status Information Alert Notices ¹ : I/O Command Set specific definition.
12	Predictable Latency Event Aggregate Log Change Notices: This bit determines whether an asynchronous event notification is sent to the host when an event pending entry for an NVM Set (refer to section 5.16.1.12) has been added to the Predictable Latency Event Aggregate Log. If this bit is set to '1', then the Predictable Latency Event Aggregate Log Change event is sent to the host when this condition occurs. If this bit is cleared to '0', then the controller shall not send the Predictable Latency Event Aggregate Log Change event to the host.
11	Asymmetric Namespace Access Change Notices: This bit determines whether an asynchronous event notification is sent to the host when an asymmetric namespace access change occurs (i.e., the contents of the Asymmetric Namespace Access log page (refer to section 5.16.1.13) change). If this bit is set to '1', then the Asymmetric Namespace Access Change Notices event is sent to the host when this condition occurs. If this bit is cleared to '0', then the controller shall not send the Asymmetric Namespace Access Change Notices event to the host.
10	Telemetry Log Notices: This bit determines whether an asynchronous event notification is sent to the host when the Telemetry Controller-Initiated Data Available field transitions from 0h to 1h in the Telemetry Controller-Initiated log page. If this bit is set to '1', then the Telemetry Log Changed event is sent to the host when this condition occurs. If this bit is cleared to '0', then the controller shall not send the Telemetry Log Changed event to the host.
09	Firmware Activation Notices: This bit determines whether an asynchronous event notification is sent to the host for a Firmware Activation Starting event (refer to Figure 146). If this bit is set to '1', then the Firmware Activation Starting event is sent to the host when this condition occurs. If this bit is cleared to '0', then the controller shall not send the Firmware Activation Starting event to the host.
08	Namespace Attribute Notices: This bit determines whether an asynchronous event notification is sent to the host for a Namespace Attribute change (refer to Figure 146). If this bit is set to '1', then the Namespace Attribute Changed event is sent to the host when this condition occurs. If this bit is cleared to '0', then the controller shall not send the Namespace Attribute Changed event to the host.

Figure 326: Asynchronous Event Configuration – Command Dword 11

Bits	Description
07:00	SMART / Health Critical Warnings: This field determines whether an asynchronous event notification is sent to the host for the corresponding Critical Warning specified in the SMART / Health Information log (refer to Figure 207). If a bit is set to '1' then an asynchronous event notification is sent when the corresponding critical warning bit is set to '1' in the SMART / Health Information log. If a bit is cleared to '0', then an asynchronous event notification is not sent when the corresponding critical warning bit is set to '1' in the SMART / Health Information log.
NOTE: 1. Refer to the NVM Command Set specification. 2. Refer to the Zoned Namespace Command Set specification.	