



LEGAL NOTICE:

© **Copyright 2007 - 2018 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 - 2018 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	4050
Change Date	2019-01-22
Builds on Specification	NVM Express 1.3c
Referenced Ratified Technical Proposal(s)	TP 4018a TP 4005a

Technical Proposal Author(s)

Name	Company
Paul Suhler	Micron Technology
Ross Stenfort	Facebook
Stacey Secatch	Seagate
Mark Carlson	Toshiba Memory Corporation
David Black	DellEMC

This technical proposal defines additional parameters for the Endurance Group Information log page; these parameters are inherited from the SMART / Health Information log page. This includes critical warnings and associated asynchronous event notifications.

It includes the ability to place all namespaces (and hence NVM Sets) in an Endurance Group in read-only mode.

Revision History

Revision Date	Change Description
2018-08-26	Initial version Corrected definition of Available Spare Threshold field Modified descriptions of pre-existing fields in the Endurance Group Information log page to avoid implying that support for Endurance Groups requires support for NVM Sets.
2018-09-18	Comments from 2018-09-13 Technical WG meeting: <ul style="list-style-type: none">Removed concept of endurance groups existing without NVM sets, which had allowed the create namespace operation to specify only an ENDGID.Moved selection of endurance group critical warnings to a new feature ID, Endurance Group Event Configuration.

2018-10-09	Comments from 2018-09-20 Technical WG meeting and subsequent e-mail exchanges: <ul style="list-style-type: none"> Deleted item from description about namespace creation and ENDGID. Modified rules relating endurance group critical warning and SMART (whole drive) critical warnings. Aligned all figure numbers to 1.3c. For figures and sections added by TPs, referenced those TPs and the figure and section numbers in the TPs. Added Endurance Group Critical Warning Summary field to the SMART / Health Information log page. Added an Entry n to the Endurance Group Event Aggregate Log Page, rather than a row with ellipses.
2018-10-23	Comments from 2018-10-11 Technical WG meeting and subsequent e-mail exchanges: <ul style="list-style-type: none"> Fixed byte of Endurance Group Critical Warning Summary field. Corrected "Compare and Read commands" to "... operations". Changed wording of new Host Read / Write Commands fields to match NVM subsystem scope of this page. Clarified parts of 8.TBD3.1.
2018-10-25	Version for phase 2 exit vote. <ul style="list-style-type: none"> Wording changes suggested by Mike Allison.
2018-11-09	Changes from discussion in 2018-11-08 Technical WG meeting: <ul style="list-style-type: none"> Removed instances of "pending" that implied that it was part of the name of an event. 5.2: Added list item for Asynchronous Event Information – Notice events, which was missing. This category includes the new Endurance Group notice. Clarifications in 8.TBD3.1.
2018-11-15	Changes from discussion in 2018-11-15 Technical WG meeting: <ul style="list-style-type: none"> Removed list item change in 5.2, as ECN 006 will make this change. Other minor wording changes. This version is ready for a Phase 3 exit vote and Member Review.
2018-12-31	Integration
2019-01-22	Ratification

Incompatible Changes

None

Description for NVMe 1.4 Changes Document

Endurance Group Information Enhancements (optional)

- Added new fields to Endurance Group Information log page:
 - Available Spare;
 - Percentage Used;
 - Host Read Commands;
 - Host Write Commands;
 - Media and Data Integrity Errors; and
 - Number of Error Information Log Entries;
- Corrected definition of Available Spare Threshold field in Endurance Group Information log page;
- Added new Endurance Group Event Aggregate log page;
- Extended Asynchronous Event Configuration Feature to enable/disable asynchronous event notifications for events added to Endurance Group Event Aggregate log page;
- Added new Endurance Group Event Configuration Feature to configure reporting of critical warnings on a per-Endurance Group basis; and
- References:

Technical input submitted to the NVM Express™ Workgroup is subject to the terms of the NVM Express™ Participant's agreement. Copyright © 2014-2018 NVMe™ Corporation.

- NVMe 1.3 sections 5.2, 5.14, 5.15, 5.20, and 5.21;
- Technical Proposal 4050;
- Technical Proposal 4018a; and
- Technical Proposal 4005a.

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

Description of Specification Changes

5 Admin Command Set

...

5.2 Asynchronous Event Request command

...

Figure 49: Asynchronous Event Information – Notice

Value	Description
00h	
02h	Telemetry Log Changed: The controller has saved the controller internal state in the Telemetry Controller-Initiated log page and set the Telemetry Controller-Initiated Data Available field to 1h in that log page. To clear this event, the host issues a Get Log Page with Retain Asynchronous Event cleared to '0' for the Telemetry Controller-Initiated Log.
06h	Endurance Group Event Aggregate Log Page Change: Indicates that event entries for one or more Endurance Groups (refer to section 5.14.1.TBD) have been added to the Predictable Latency Event Aggregate log. To clear this event, the host issues a Get Log Page command with Retain Asynchronous Event bit cleared to '0' for the Endurance Group Event Aggregate log.
07h to FFh	Reserved

<Editor's Note: The above TBD bit number is assigned when the TP is sent for ratification. That will resolve a race condition with other TPs requesting new values in this figure.>

<Editor's Note: Section 5.14.1.TBD was introduced in TP 4018a.>

...

5.14 Get Log Page Command

...

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 and the scope of the information that is returned in those Log pages.

Log pages that indicate a scope of NVM subsystem return information that is global to the NVM subsystem. Log pages that indicate a scope of controller return information that is specific to the controller that is processing the command. Log pages that indicate a scope of Namespace return information that is specific to the specified namespace. For log pages that indicate multiple scopes, the namespace identifier that is specified determines which information is returned. The definition of any individual field within a Log page may indicate a different scope that is specific to that individual field.

For Log Pages with a scope of NVM subsystem or controller (as shown in Figure 90 and Figure 91), the controller should abort commands that specify namespace identifiers other than 0h or FFFFFFFFh with status Invalid Field in Command. Otherwise the rules for namespace identifier usage in Error! Reference source not found. apply.

Figure 90: Get Log Page – Log Page Identifiers

Log Identifier	O/M	Scope	Description	Reference Section
00h	Reserved			
01h	M	Controller	Error Information	Error! Reference source not found.
02h	M	Controller ¹	SMART / Health Information	Error! Reference source not found.
	O	Namespace ²		

Figure 90: Get Log Page – Log Page Identifiers

Log Identifier	O/M	Scope	Description	Reference Section
03h	M	NVM subsystem	Firmware Slot Information	Error! Reference source not found.
04h	O	Controller	Changed Namespace List	Error! Reference source not found.
05h	O	Controller	Commands Supported and Effects	Error! Reference source not found.
06h	O	Controller ³	Device Self-test	Error! Reference source not found.
	O	NVM subsystem ⁴		
07h	O	Controller	Telemetry Host-Initiated	Error! Reference source not found.
08h	O	Controller	Telemetry Controller-Initiated	Error! Reference source not found.
0Fh	O	NVM subsystem	Endurance Group Event Aggregate	5.14.1.TBD
09h to 6Fh	Reserved			
70h	Discovery (refer to the NVMe over Fabrics specification)			
71h to 7Fh	Reserved for NVMe over Fabrics implementations			
80h to BFh	I/O Command Set Specific			
C0h to FFh	Vendor specific			
KEY: O/M definition: 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 Figure 2). 4. Bit 0 is set to '1' in the DSTO field in the Identify Controller data structure.				

<Editor's Note: The above TBD value is assigned when the TP is sent for ratification. That will resolve a race condition with other TPs requesting new log pages.>

<Editor's Note: Section 5.14.1.TBD was introduced in TP 4018a.>

...

The following change to the Endurance Group Information log page is a change to TP 4018a:

5.14.1.2 SMART / Health Information (Log Identifier 02h)

...

The log page returned is defined in Figure 1.

Figure 1: Get Log Page – SMART / Health Information Log

Bytes	Description														
00	<p>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. 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.</p> <table> <tr> <th>Bit</th><th>Definition</th></tr> <tr> <td>0</td><td>If set to '1', then the available spare capacity has fallen below the threshold.</td></tr> <tr> <td>1</td><td>If set to '1', then a temperature is above an over temperature threshold or below an under temperature threshold (refer to section Error! Reference source not found.).</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>3</td><td>If set to '1', then the media has been placed in read only mode.</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>7:5</td><td>Reserved</td></tr> </table>	Bit	Definition	0	If set to '1', then the available spare capacity has fallen below the threshold.	1	If set to '1', then a temperature is above an over temperature threshold or below an under temperature threshold (refer to section Error! Reference source not found.).	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.	3	If set to '1', then the media has been placed in read only mode.	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.	7:5	Reserved
Bit	Definition														
0	If set to '1', then the available spare capacity has fallen below the threshold.														
1	If set to '1', then a temperature is above an over temperature threshold or below an under temperature threshold (refer to section Error! Reference source not found.).														
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.														
3	If set to '1', then the media has been placed in read only mode.														
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.														
7:5	Reserved														
02:01	<p>Composite Temperature: Contains a value corresponding to a temperature in degrees Kelvin that represents the current composite temperature of the controller and namespace(s) associated with that controller. The 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 Error! Reference source not found.).</p> <p>Warning and critical overheating composite temperature threshold values are reported by the WCTEMP and CCTEMP fields in the Identify Controller data structure in Error! Reference source not found.</p>														
03	Available Spare: Contains a normalized percentage (0 to 100%) of the remaining spare capacity available.														
04	Available Spare Threshold: When the Available Spare falls below the threshold indicated in this field, an asynchronous event completion may occur. The value is indicated as a normalized percentage (0% to 100%). The values 101 to 255 are reserved.														
05	<p>Percentage Used: Contains a vendor specific estimate of the percentage of NVM subsystem life used based on the actual usage and the manufacturer's prediction of NVM life. A value of 100 indicates that the estimated endurance of the NVM in the NVM subsystem has been consumed, but may not indicate an NVM subsystem failure. The value is allowed to exceed 100. Percentages greater than 254 shall be represented as 255. This value shall be updated once per power-on hour (when the controller is not in a sleep state).</p> <p>Refer to the JEDEC JESD218A standard for SSD device life and endurance measurement techniques.</p>														

Bytes	Description												
06	<p>Endurance Group Critical Warning Summary: This field indicates critical warnings for the state of Endurance Groups. Each bit corresponds to a critical warning type, multiple bits may be set. If a bit is cleared to '0', then that critical warning does not apply to any Endurance Group. 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.</p> <p>If a bit is set to '1' in one or more Endurance Groups, then the corresponding bit shall be set to '1' in this field.</p> <table> <tr> <th>Bit</th><th>Definition</th></tr> <tr> <td>0</td><td>If set to '1', then the available spare capacity of one or more Endurance Groups has fallen below the threshold.</td></tr> <tr> <td>1</td><td>Reserved</td></tr> <tr> <td>2</td><td>If set to '1', then the reliability of one or more Endurance Groups has been degraded due to significant media related errors or any internal error that degrades NVM subsystem reliability.</td></tr> <tr> <td>3</td><td>If set to '1', then the namespaces in one or more Endurance Groups have been placed in read only mode not as a result of a change in the write protection state of a namespace (refer to section 8.TBD).</td></tr> <tr> <td>7:4</td><td>Reserved</td></tr> </table>	Bit	Definition	0	If set to '1', then the available spare capacity of one or more Endurance Groups has fallen below the threshold.	1	Reserved	2	If set to '1', then the reliability of one or more Endurance Groups has been degraded due to significant media related errors or any internal error that degrades NVM subsystem reliability.	3	If set to '1', then the namespaces in one or more Endurance Groups have been placed in read only mode not as a result of a change in the write protection state of a namespace (refer to section 8.TBD).	7:4	Reserved
Bit	Definition												
0	If set to '1', then the available spare capacity of one or more Endurance Groups has fallen below the threshold.												
1	Reserved												
2	If set to '1', then the reliability of one or more Endurance Groups has been degraded due to significant media related errors or any internal error that degrades NVM subsystem reliability.												
3	If set to '1', then the namespaces in one or more Endurance Groups have been placed in read only mode not as a result of a change in the write protection state of a namespace (refer to section 8.TBD).												
7:4	Reserved												
631:07	Reserved												
47:32	<p>Data Units Read: Contains the number of 512 byte data units the host has read from the controller; this value does not include metadata. This value is reported in thousands (i.e., a value of 1 corresponds to 1000 units of 512 bytes read) and is rounded up. When the LBA size is a value other than 512 bytes, the controller shall convert the amount of data read to 512 byte units.</p> <p>For the NVM command set, logical blocks read as part of Compare and Read operations shall be included in this value.</p>												
...													

<Editor's Note: The above TBD value is assigned when the TP is sent for ratification. That will resolve a race condition with other TPs requesting new fields in this log page.>

< Editor's Note: Section 8.TBD was introduced in TP 4005a. >

...

5.14.1.9 Endurance Group Information (Log Identifier 09h)

This log page is used to provide endurance information based on the Endurance Group (refer to [section TBD.3](#)). An Endurance Group consists of zero or more NVM Sets. The information provided is over the life of the Endurance Group. The Endurance Group Identifier is specified in the Log Specific Identifier field in Command Dword 11 of the Get Log Page command. The log page is 512 bytes in size.

Figure TP_4018a_5_14_1_11Fig0: Get Log Page – Endurance Group Information (Log Identifier 09h)

Bytes	Description												
00	<p>Critical Warning: This field indicates critical warnings for the state of the Endurance Group. Each bit corresponds to a critical warning type; multiple bits may be set. 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.</p> <p>If a bit is set to '1' in all Endurance Groups in the NVM subsystem, then the corresponding bit shall be set to '1' in the Endurance Group Critical Warnings field of the SMART / Health Information log page (refer to Figure 93).</p> <table border="1"> <thead> <tr> <th>Bit</th><th>Definition</th></tr> </thead> <tbody> <tr> <td>0</td><td>If set to '1', then the available spare capacity of the Endurance Group has fallen below the threshold.</td></tr> <tr> <td>1</td><td>Reserved</td></tr> <tr> <td>2</td><td>If set to '1', then the Endurance Group reliability has been degraded due to significant media related errors or any internal error that degrades NVM subsystem reliability.</td></tr> <tr> <td>3</td><td>If set to '1', then all namespaces in the Endurance Group have been placed in read only mode for reasons other than a change in the write protect state of the namespace. The controller shall not set this bit to '1' if the read-only condition on the Endurance Group is a result of a change in the write protection state of all namespaces in the Endurance Group.</td></tr> <tr> <td>7:4</td><td>Reserved</td></tr> </tbody> </table>	Bit	Definition	0	If set to '1', then the available spare capacity of the Endurance Group has fallen below the threshold.	1	Reserved	2	If set to '1', then the Endurance Group reliability has been degraded due to significant media related errors or any internal error that degrades NVM subsystem reliability.	3	If set to '1', then all namespaces in the Endurance Group have been placed in read only mode for reasons other than a change in the write protect state of the namespace. The controller shall not set this bit to '1' if the read-only condition on the Endurance Group is a result of a change in the write protection state of all namespaces in the Endurance Group.	7:4	Reserved
Bit	Definition												
0	If set to '1', then the available spare capacity of the Endurance Group has fallen below the threshold.												
1	Reserved												
2	If set to '1', then the Endurance Group reliability has been degraded due to significant media related errors or any internal error that degrades NVM subsystem reliability.												
3	If set to '1', then all namespaces in the Endurance Group have been placed in read only mode for reasons other than a change in the write protect state of the namespace. The controller shall not set this bit to '1' if the read-only condition on the Endurance Group is a result of a change in the write protection state of all namespaces in the Endurance Group.												
7:4	Reserved												
3:0 02:01	Reserved												
03	Available Spare: Contains a normalized percentage (0% to 100%) of the remaining spare capacity available for the Endurance Group.												
04	Available Spare Threshold: The available spare If the Available Spare falls below the threshold indicated in this field, an asynchronous event completion may occur. The value is indicated as a normalized percentage (0% to 100%). The values 101 to 255 are reserved.												
05	Percentage Used: Contains a vendor specific estimate of the percentage of life used for the NVM Set(s) that comprise the Endurance Group based on the actual usage and the manufacturer's prediction of NVM life. A value of 100 indicates that the estimated endurance of the NVM in the Endurance Group has been consumed, but may not indicate an NVM failure. The value is allowed to exceed 100. Percentages greater than 254 shall be represented as 255. This value shall be updated once per power-on hour when the controller is not in a sleep state. Refer to the JEDEC JESD218A standard for SSD device life and endurance measurement techniques.												
31:06	Reserved												
47:32	Endurance Estimate: This field is an estimate of the total number of data bytes that may be written to the NVM Set(s) that comprise the Endurance Group over the lifetime of the Endurance Group assuming a Write Amplification of 1. This value is reported in billions (i.e., a value of 1 corresponds to 1,000,000,000 bytes written) and is rounded up. A value of zero indicates that the controller does not report an Endurance Estimate.												
63:48	Data Units Read: Contains the total number of data bytes that have been read from the NVM Set(s) that comprise the Endurance Group. This value does not include controller reads due to internal operations such as garbage collection. This value is reported in billions (i.e., a value of 1 corresponds to 1,000,000,000 bytes read) and is rounded up. A value of zero indicates that the controller does not report the number of Data Units Read.												

Figure TP_4018a_5_14_1_11Fig0: Get Log Page – Endurance Group Information (Log Identifier 09h)

Bytes	Description
79:64	Data Units Written: Contains the total number of data bytes that have been written to the NVM Set(s) that comprise the Endurance Group. This value does not include controller writes due to internal operations such as garbage collection. This value is reported in billions (i.e., a value of 1 corresponds to 1,000,000,000 bytes written) and is rounded up. A value of zero indicates that the controller does not report the number of Data Units Written.
95:80	Media Units Written: Contains the total number of data bytes that have been written to the NVM Set(s) that comprise the Endurance Group including both host and controller writes (e.g., garbage collection). This value is reported in billions (i.e., a value of 1 corresponds to 1,000,000,000 bytes written) and is rounded up. A value of zero indicates that controller does not report the number of Media Units Written.
111:96	Host Read Commands: Contains the number of read commands completed by all controllers in the NVM subsystem for the Endurance Group. For the NVM command set, this is the number of Compare commands and Read commands.
127:112	Host Write Commands: Contains the number of write commands completed by all controllers in the NVM subsystem for the Endurance Group. For the NVM command set, this is the number of Write commands.
143:128	Media and Data Integrity Errors: Contains the number of occurrences where the controller detected an unrecovered data integrity error for the Endurance Group. Errors such as uncorrectable ECC, CRC checksum failure, or LBA tag mismatch are included in this field.
159:144	Number of Error Information Log Entries: Contains the number of Error Information log entries over the life of the controller for the Endurance Group.
511:96 160	Reserved

< Editor's Note: Section 8.TBD3 was introduced in TP 4018a. >

< Editor's Note: The above figure TP_4018a_5_14_1_11Fig0 was introduced in TP 4018a. >

< Editor's Note: Section 8.TBD2.1 was introduced in TP 4005a. >

...

5.14.1.TBD Endurance Group Event Aggregate (Log Identifier 0Fh)

This log page indicates if an Endurance Group Event (refer to section 0) has occurred for a particular Endurance Group. If an Endurance Group Event has occurred, the details of the particular event are included in the Endurance Group Information log page for that Endurance Group. An asynchronous event is generated when an entry for an Endurance Group is newly added to this log page.

If there is an enabled Endurance Group Event pending for an Endurance Group, then the Endurance Group Event Aggregate log page includes an entry for that Endurance Group. The log page is an ordered list by Endurance Group Identifier. For example, if Endurance Group Events are pending for Endurance Group 2, 1, and 7, then the log page shall have entries in numerical order of 1, 2, and 7. A particular Endurance Group entry is removed from this log page after the Get Log Page is completed successfully with the Retain Asynchronous Event bit cleared to '0' for the Endurance Group Information log page for that Endurance Group.

The log page size is limited by the Endurance Group Identifier Maximum value reported in the Identify Controller data structure (refer to Figure 111). If the host reads beyond the end of the log page, zeros are returned. The log page is defined in Figure .

Figure TP4050_FigTBD: Get Log Page – Endurance Group Event Aggregate Log Page

Bytes	Description
07:00	Number of Entries: This field indicates the number of entries in the list. The maximum number of entries in the list corresponds to the Endurance Group Identifier Maximum field reported in the Identify Controller data structure. A value of 0h indicates there are no entries in the list.
09:08	Entry 1: Indicates the Endurance Group that has an Endurance Group Event pending that has the numerically smallest Endurance Group Identifier, if any.
11:10	Entry 2: Indicates the Endurance Group that has an Endurance Group Event pending that has the second numerically smallest Endurance Group Identifier, if any.
13:12	Entry 3: Indicates the Endurance Group that has an Endurance Group Event pending that has the third numerically smallest Endurance Group Identifier, if any.
15:14	Entry 4: Indicates the Endurance Group that has an Endurance Group Event pending that has the fourth numerically smallest Endurance Group Identifier, if any.
...	...
2*n+7: 2*n+6	Entry n: Indicates the Endurance Group that has an Endurance Group Event pending that has the numerically largest Endurance Group Identifier, if any.

< Editor's Note: Section 8.TBD3 was introduced in TP 4018a. >

5.15 Identify command

5.15.1 Identify command overview

...

5.15.2 Identify Data Structures

...

5.15.2.2 Identify Controller data structure (CNS 01h)

The Identify Controller data structure (refer to **Figure 2**) is returned to the host for this controller.

Figure 2: Identify – Identify Controller Data Structure

Bytes	O/M ¹	Description
Controller Capabilities and Features		
...		
95:92	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>Bits 31:15 are reserved.</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 9 is set to '1' if the controller supports sending Firmware Activation Notices. 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 sending Namespace Attribute Notices 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>

Figure 2: Identify – Identify Controller Data Structure

Bytes	O/M ¹	Description
...		
339:338	O	NVM Set Identifier Maximum (NSETIDMAX): This field defines the maximum value of a valid NVM Set Identifier for any controller in the NVM subsystem. The number of NVM Sets supported by the NVM subsystem is less than or equal to NSETIDMAX.
341:340	O	Endurance Group Identifier Maximum (ENDGIDMAX): This field defines the maximum value of a valid Endurance Group Identifier for any controller in the NVM subsystem. The number of Endurance Groups supported by the NVM subsystem is less than or equal to ENDGIDMAX.
...		
NOTES: 1 O/M definition: O = Optional, M = Mandatory.		

<Editor's Note: The above NVMSETIDMAX field was added in TP 4018a. It is recommended to add the ENDGIDMAX field immediately following that, if possible. >

...

5.21 Set Features command

...

5.21.1 Feature Specific Information

...

Figure 3: Set Features – Feature Identifiers

Feature Identifier	O/M ⁶	Persistent Across Power Cycle and Reset ²	Uses Memory Buffer for Attributes	Description
00h				Reserved
01h	M	No	No	Arbitration
02h	M	No	No	Power Management
03h	O	Yes	Yes	LBA Range Type
04h	M	No	No	Temperature Threshold
05h	M	No	No	Error Recovery
06h	O	No	No	Volatile Write Cache
07h	M	No	No	Number of Queues
08h	NOTE 5	No	No	Interrupt Coalescing
09h	NOTE 5	No	No	Interrupt Vector Configuration
0Ah	M	No	No	Write Atomicity Normal
0Bh	M	No	No	Asynchronous Event Configuration
0Ch	O	No	Yes	Autonomous Power State Transition
0Dh	O	No ³	No ⁴	Host Memory Buffer
0Eh	O	No	Yes	Timestamp
0Fh	O	No	No	Keep Alive Timer
10h	O	Yes	No	Host Controlled Thermal Management
11h	O	No	No	Non-Operational Power State Config
18h	O	No	No	Endurance Group Event Configuration
12h 19h – 77h				Reserved
78h – 7Fh		Refer to the NVMe Management Interface Specification for definition.		
80h – BFh				Command Set Specific (Reserved)
C0h – FFh				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 **Error! Reference source not found.**). 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 **Error! Reference source not found.**
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 **Error! Reference source not found.**
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.

<Editor's Note: The above TBD bit number is assigned when the TP is sent for ratification. That will resolve a race condition with other TPs requesting new values in this figure.>

...

5.21.1.11 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 Error! Reference source not found.**). 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 indicated in Command Dword 11.

If a Get Features command is submitted for this Feature, the attributes specified in **Figure 143** are returned in Dword 0 of the completion queue entry for that command.

Figure 143: Asynchronous Event Configuration – Command Dword 11

Bits	Description
31:15	Reserved
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 8.TBD3) has been added to the Endurance Group Event Aggregate log (refer to section 5.14.1.TBD). 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.
10	Telemetry Log Notices: This field determines whether an asynchronous event notification is sent to the host when the Telemetry Controller-Initiated Data Available field transitions from '0' to '1' 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 field determines whether an asynchronous event notification is sent to the host for a Firmware Activation Starting event (refer to Figure). 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 field determines whether an asynchronous event notification is sent to the host for a Namespace Attribute change (refer to Figure). 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.
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 Error! Reference source not found.). 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.

< Editor's Note: Section 8.TBD3 was introduced in TP 4018a. >

<Editor's Note: The above TBD bit number is assigned when the TP is sent for ratification. That will resolve a race condition with other TPs requesting new values in this figure.>

...

5.21.1.TBD Endurance Group Event Configuration (Feature Identifier 18h), (Optional)

This Feature controls the events that trigger adding an Endurance Group Event Aggregate Log Change Notices event to the Endurance Group Event Aggregate log. 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 indicated in Command Dword 11.

If a Get Features command is submitted for this Feature, the Endurance Group Critical Warnings field in Command Dword 11 is not used and the attributes specified in **Figure TP4050_FigTBD2** are returned in Dword 0 of the completion queue entry for that command.

Figure TP4050_FigTBD2: Asynchronous Event Configuration – Command Dword 11

Bits	Description
31:24	Reserved
23:16	Endurance Group Critical Warnings: This field determines whether an event entry for an Endurance Group (refer to section 8.TBD3) is added to the Endurance Group Event Aggregate log page (refer to section 5.14.1.TBD) for the corresponding Critical Warning specified in the Endurance Group Information log (refer to Figure TP_4018a_5_14_1_11Fig0). If a bit is set to '1', then an entry is added when the corresponding critical warning bit is set to '1' in the Endurance Group Information log. If a bit is cleared to '0', then an entry is not added when the corresponding critical warning bit is set to '1' in the Endurance Group Information log page.
15:00	Endurance Group Identifier (ENDGID): This field indicates the Endurance Group for which asynchronous events are being configured. If this field is cleared to 0h, then the Endurance Group Critical Warnings field is not used.

If a bit is set to '1' in the Endurance Group Critical Warnings field which corresponds to a reserved bit in the Critical Warning field of the Endurance Group Information log page (refer to Figure 5_14_1_11Fig0), then the Set Features command shall be aborted with a status of Invalid Field in Command.

If the Endurance Group Identifier specifies an Endurance Group that does not exist, then the Set Features or Get Features command shall be aborted with a status of Invalid Field in Command.

< Editor's Note: Section 8.TBD3 was introduced in TP 4018a. >

< Editor's Note: Figure 5_14_1_11Fig0 was introduced in TP 4018a. >

5.21.2 Command Completion

...

8 Features

...

The following is a change to TP 4018a:

< Editor's Note: Section 8.TBD3 was introduced in TP 4018a. The text below is added at the end of that section. >

8.TBD3 Endurance Groups

...

8.TBD3.1 Configuring and Managing Events

The host may configure asynchronous events to be triggered when certain events occur for an Endurance Group. The host submits a Set Features command specifying the Endurance Group Event Configuration feature (refer to section 5.21.1.TBD), the Endurance Group, and the specific event(s) that shall trigger adding an entry to the Endurance Group Event Aggregate log page (refer to section 5.14.1.TBD).

The host configures events using a Set Feature command for each Endurance Group.

The host submits a Set Features command specifying the Asynchronous Event Configuration feature (refer to section 5.21.1.11) with the Endurance Group Event Aggregate Log Change Notices bit set to '1' to specify that adding an entry to the Endurance Group Event Aggregate log page shall trigger an Endurance Group Event Aggregate Log Page Change Notice event to the host (refer to Figure TP4050_FigTBD2).

The host determines the Endurance Groups that have outstanding events by reading the Endurance Group Event Aggregate log page. An entry is returned for each Endurance Group that has an event outstanding. The host may use the Endurance Group Identifier Maximum value reported in the Identify Controller data structure to determine the maximum size of this log page.

To determine the specific event(s) that have occurred for a reported Endurance Group, the host reads the Endurance Group log page for that Endurance Group. The Critical Warning field indicates the event(s) that

have occurred (e.g., that all namespaces in the Endurance Group have been placed in read-only mode). All events for an Endurance Group are cleared if the controller successfully processes a read for the Endurance Group Information log page for that Endurance Group, where the Get Log Page command has the Retain Asynchronous Event bit cleared to '0'. If the Critical Warning field in the Endurance Group Information log page is cleared to 0h, then events for that Endurance Group are not reported in the Endurance Group Event Aggregate log page.

<End of changes>