



LEGAL NOTICE:

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

This NVM Express revision 1.4 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 1.4 technical proposal subject, however, to the Member's continued compliance with the Company's Intellectual Property Policy and Bylaws and the Member's Participation Agreement.

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

LEGAL DISCLAIMER:

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

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

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

NVM Express Technical Proposal for New Feature

Technical Proposal ID	4063 Telemetry Log Enhancement
Change Date	7/22/2019
Builds on Specification	NVM Express 1.4
Ratified Technical Proposals Referenced	None

Technical Proposal Author(s)

Name	Company
Michael Allison, Lukasz Tur, Randal Eike	Intel Technology

This technical proposal allows a controller to indicate to the host if the Controller-Initiated log page is allowed to be overwritten by the controller.

Revision History

Revision Date	Change Description
2019-04-01	Initial version
2019-04-03	Added clarity to as to when the RAE bit is to be cleared to '0' or set to '1' for the Controller-Initiated log.
2019-04-11	Added the word "Telemetry" where needed. Added a vendor specific time to allow the controller to remove old entries of the controller-initiated log.
2019-05-06	Aligned host-initiated logs to have a generation number for multiple controller support. Changed the scope of the host-initiated and controller-initiated logs to be vendor specific. Removed the bit for the controller to declare locking the controller-initiated log until full read by host.
2019-05-31	Reworded the rolls over text to be consistent with editorial changes during NVMe 1.4 draft review. Removed the word "it" to align with NVMe 1.4 draft review. Removed excess text that was not changing (most to avoid editorial changes in final version of NVMe 1.4).
2019-06-06	Removed tracking changes for 30 day member review. Removed a repeated word.
2019-06-14	Editorial changes from Edward Hsieh from member review.
2019-07-12	Integrated
2019-07-18	Technical WG adjustments on editorial items. Approved by WG for ratification.
2019-07-22	Ratified

Description for NVMe 1.4+ Changes Document

Allow controllers to report if the Controller-Initiated log page is not overwritten (optional)

Technical input submitted to the NVM Express™ Workgroup is subject to the terms of the NVM Express™ Participant's agreement. Copyright © 2014-2019 NVMe™ Corporation.

- Add a generation number in the Host-Initiated log page
- Aligned the process of reading the Host-Initiated log page and the Controller-Initiated log page
- Added the ability for Host-Initiated log page to have either controller scope of NVM subsystem scope.

Specification Changes

Markup Conventions:

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

Modify a portion of Figure 191 (Get Log Page – Log Page Identifiers) as shown below:

Figure 1: Get Log Page – Log Page Identifiers

Log Identifier	Scope	Log Page Name	Reference Section
...			
07h	Controller Vendor Specific	Telemetry Host-Initiated ⁵	5.14.1.7
08h	Controller Vendor Specific	Telemetry Controller-Initiated ⁵	5.14.1.8
KEY: 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. Vendor Specific = The log page contains information that is vendor specific.			

Modify a portion of Figure 202 (Get Log Page – Telemetry Host-Initiated Log (Log Identifier 07h)) as shown below:

Figure 2: Get Log Page – Telemetry Host-Initiated Log (Log Identifier 07h)

Bytes	Description
...	
38 40 :14	Reserved
381	Telemetry Host-Initiated Data Generation Number: Contains a value that is incremented each time the controller captures its internal controller state for this log page. If the value of this field is FFh, then the field shall be cleared to 0h when incremented (i.e., rolls over to 0h).

Modify a portion of section 8.14 (Telemetry (Optional)) as shown below:

...

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

1. The host identifies controller support for Telemetry log pages in the Identify Controller data structure;
2. The host prepares an area to store telemetry data if needed;
3. To receive notification that controller-initiated telemetry data is available, the host enables Telemetry Log Notices using the Asynchronous Event Configuration feature (refer to section 5.21.1.11); and
4. If the host decides to collect host-initiated telemetry data or the controller signals that controller-initiated telemetry data is available:
 - a. The host reads the appropriate blocks of the Telemetry Data Area from the host-initiated log (refer to section 5.14.1.7) or the controller-initiated log (refer to section 5.14.1.8). If possible, the host should collect Telemetry Data Area 1, 2, and 3. The host reads the log in 512 byte Telemetry Data Block units. ~~The host should As part of the last read for a controller-initiated log, the host clears set~~ the Retain Asynchronous Event bit to '01';
 - b. ~~If it is a controller-initiated log, then it~~ The host re-reads the header of the log page and ensures that the ~~Telemetry Host-Initiated Data Generation Number field from the host-initiated log or the~~ Telemetry Controller-Initiated Data Generation Number field in the ~~controller-initiated log~~ matches the original value read. If these values do not match, then the data captured is not consistent and should be re-read from the ~~log page with the Retain Asynchronous Event bit set to '1' ;~~;
 - c. ~~If the host is reading the controller-initiated log, then the host reads any portion of that log page with the Retain Asynchronous Event bit cleared to '0' to indicate to the controller that the host has completed reading the controller-initiated log page;~~ and
 - d. When all telemetry data has been saved, the data should be forwarded to the manufacturer of the controller.

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

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

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