



LEGAL NOTICE:

© Copyright 2007 - 2015 NVM Express, Inc. ALL RIGHTS RESERVED.

This NVM Express revision 1.2 specification 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 1.2 specification 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 - 2015 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 Virtual, Inc.
401 Edgewater Place, Suite 600
Wakefield, MA 01880
info@nvmexpress.org

NVM Express™ Technical Errata

Errata ID	006
Revision Date	11/18/2015
Affected Spec Ver.	NVM Express™ 1.2
Corrected Spec Ver.	

Errata Author(s)

Name	Company
John Carroll	Intel
Matthew Wilcox	Intel
Judy Brock	Samsung
Jim Hatfield	Seagate
Dave Landsman	SanDisk
Nadesan Narenthiran	HGST
Ken Okin	KOC

Errata Overview

Updates include:

- Clarification to Write Uncorrectable command specific status values.
- Specifying that NGUID and EUI64 include the Organizationally Unique Identifier.
- Clarifications for SGL Bit Bucket.
- Added new reference for Security Features for SCSI Commands.
- Modifications to the error log that are recommendations that align with requirements in NVMe 1.2 Technical Proposal 005.

Revision History

Revision Date	Change Description
5/28/2015	First draft
07/26/2015	Added Read-Only behavior clarifications
08/20/2015	Added NGUID, Bit Bucket, and Namespace Attachment clarifications. Added new reference, and removed error log clarifications for ECN 007.
8/25/2015	Clarified Firmware Slot Information log. Removed duplicate PRPs. Changed Metadata Settings abbreviation to avoid conflict.
9/3/2015	Added Error Log items back. Modified Length for bit buckets based on reflector discussion. Moved read only error to ECN 007.
9/4/2015	Accepted changes and modified overview before sending to Board for 30-day review approval.
11/18/2015	Ratified.

Description of Specification Changes

Modify a portion of section 3.1.5 (Controller Configuration) as shown below:

00	RW	0	<p>Enable (EN): When set to '1', then the controller shall process commands based on Submission Queue Tail doorbell writes. When cleared to '0', then the controller shall not process commands nor post completion queue entries to Completion Queues. When this field transitions from '1' to '0', the controller is reset (referred to as a Controller Reset). The reset deletes all I/O Submission Queues and I/O Completion Queues, resets the Admin Submission Queue and Completion Queue, and brings the hardware to an idle state. The reset does not affect PCI Express registers nor the Admin Queue registers (AQA, ASQ, or ACQ). All other controller registers defined in this section and internal controller state (e.g., Feature values defined in section 5.14.1 that are not persistent across power states) are reset to their default values. The controller shall ensure that there is no data loss for commands that have had corresponding completion queue entries posted to an I/O Completion Queue prior to the reset operation. Refer to section Error! Reference source not found. for reset details.</p> <p>When this field is cleared to '0', the CSTS.RDY bit is cleared to '0' by the controller once the controller is ready to be re-enabled. When this field is set to '1', the controller sets CSTS.RDY to '1' when it is ready to process commands. CSTS.RDY may be set to '1' before namespace(s) are ready to be accessed.</p> <p>Setting this field from a '0' to a '1' when CSTS.RDY is a '1,' or setting this field from a '1' to a '0' when CSTS.RDY is a '0,' has undefined results. The Admin Queue registers (AQA, ASQ, and ACQ) shall only be modified when EN is cleared to '0'.</p>
----	----	---	--

Modify a portion of Figure 79 (SMART Log Page) as shown below:

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 the Critical Composite Temperature Threshold (CCTEMP) field in the Identify Controller data structure in Error! Reference source not found.</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>
---------	--

Modify a portion of section 5.12 (Namespace Attachment Command) as shown below:

The Namespace Attachment command uses the **PRP Entry 1, PRP Entry 2, and Command Dword 10** field. All other command specific fields are reserved.

Modify a portion of Figure 108 (Set Features – Feature Identifiers) as shown below:

Feature Identifier	O/M	Persistent Across Power States Cycle and Reset ²	Uses Memory Buffer for Attributes	Description
--------------------	-----	---	-----------------------------------	-------------

Modify a portion of section 5.7.1 (Firmware Commit – Command Completion) as shown below:

For requests that specify activation of a new firmware image **at the next reset** and return with status code value of 00h, any Controller Level Reset defined in section **Error! Reference source not found.** activates the specified firmware.

Add the following figure after section 6.15.1 (Command Completion – Write Uncorrectable command)

Figure TBD: Write Uncorrectable – Command Specific Status Values

Value	Description
82h	Attempted Write to Read Only Range: The LBA range specified contains read-only blocks.

Modify a portion of Figure 92 (Identify – Identify Namespace Data Structure) as shown below:

119:104	O	<p>Namespace Globally Unique Identifier (NGUID): This field contains a 128-bit value that is globally unique and assigned to the namespace when the namespace is created. This field remains fixed throughout the life of the namespace and is preserved across namespace and controller operations (e.g., controller reset, namespace format, etc.).</p> <p>This field uses the EUI-64 based 16-byte designator format. Bytes 114:112 contain the 24-bit company_id Organizationally Unique Identifier (OUI) value assigned by the IEEE Registration Authority. Bytes 119:115 contain an extension identifier assigned by the corresponding organization. Bytes 111:104 contain the vendor specific extension identifier assigned by the corresponding organization. See the IEEE EUI-64 guidelines for more information.</p> <p>The controller shall specify a globally unique namespace identifier in this field or the EUI64 field when the namespace is created.</p>
127:120	O	<p>IEEE Extended Unique Identifier (EUI64): This field contains a 64-bit IEEE Extended Unique Identifier (EUI-64) that is globally unique and assigned to the namespace when the namespace is created. This field remains fixed throughout the life of the namespace and is preserved across namespace and controller operations (e.g., controller reset, namespace format, etc.).</p> <p>The EUI-64 is a concatenation of a 24-bit or 36-bit company_id Organizationally Unique Identifier (OUI or OUI-36) value assigned by the IEEE Registration Authority and an extension identifier assigned by the corresponding organization. See the IEEE EUI-64 guidelines for more information.</p> <p>The controller shall specify a globally unique namespace identifier in this field or the NGUID field when the namespace is created. If the controller is not able to allocate a globally unique 64-bit identifier then this field shall be cleared to 0h. Refer to section 7.9.</p>

Modify a portion of Figure 21 (SGL Bit Bucket descriptor) as shown below:

Bytes	Description
11:8	<p>Length: The Length field specifies the amount of source data that is discarded. An SGL Bit Bucket descriptor specifying that no source data be discarded (i.e., the length field set to 00000000h) is a valid SGL Bit Bucket descriptor.</p> <p>If the SGL Bit Bucket Descriptor describes a destination data buffer (e.g., a read from the controller to memory), then the Length field specifies the number of bytes of the source data which the controller shall to be discarded (i.e., not transfered to the destination data buffer).</p> <p>If the SGL Bit Bucket Descriptor describes a source data buffer (e.g., a write from memory to the controller), then the Length field the Bit Bucket descriptor shall be ignored-treated as if the Length field were set to 00000000h (i.e., the Bit Bucket Descriptor has no effect).</p> <p>If SGL Bit Bucket descriptors are supported, their length in a destination data buffer shall be included in the Number of Logical Blocks (NLB) parameter specified in NVM Command Set data transfer commands. Their length in a source data buffer is not included in the NLB parameter.</p>

Modify a portion of section 1.9 (References) as shown below:

Eastlake, D. and T. Hansen, "US Secure Hash Algorithms (SHA and 4318 HMAC-SHA)", RFC 4634, July 2006.

INCITS 501-201x, Information technology – Security Features for SCSI Commands (SFSC)

PCI specification, revision 3.0. Available from <http://www.pcisig.com>.

Modify a portion of Figure 81 (Firmware Slot Information Log) as shown below:

Figure 1: Get Log Page – Firmware Slot Information Log

Bytes	Description
00	<p>Active Firmware Info (AFI): Specifies information about the active firmware revision.</p> <p>Bit 7 is reserved.</p> <p>Bits 6:4 indicates the firmware slot that is going to be activated at the next controller reset. If this field is 0h, then the controller does not indicate the firmware slot that is going to be activated at the next controller reset.</p> <p>Bit 3 is reserved.</p> <p>Bits 2:0 indicates the firmware slot from which the actively running firmware revision was loaded.</p>
07:01	Reserved

15:08	Firmware Revision for Slot 1 (FRS1): Contains the revision of the firmware downloaded to firmware slot 1. If no valid firmware revision is present or if this slot is unsupported, all zeros shall be returned this field shall be cleared to 0h.
23:16	Firmware Revision for Slot 2 (FRS2): Contains the revision of the firmware downloaded to firmware slot 2. If no valid firmware revision is present or if this slot is unsupported, all zeros shall be returned this field shall be cleared to 0h.
31:24	Firmware Revision for Slot 3 (FRS3): Contains the revision of the firmware downloaded to firmware slot 3. If no valid firmware revision is present or if this slot is unsupported, all zeros shall be returned this field shall be cleared to 0h.
39:32	Firmware Revision for Slot 4 (FRS4): Contains the revision of the firmware downloaded to firmware slot 4. If no valid firmware revision is present or if this slot is unsupported, all zeros shall be returned this field shall be cleared to 0h.
47:40	Firmware Revision for Slot 5 (FRS5): Contains the revision of the firmware downloaded to firmware slot 5. If no valid firmware revision is present or if this slot is unsupported, all zeros shall be returned this field shall be cleared to 0h.
55:48	Firmware Revision for Slot 6 (FRS6): Contains the revision of the firmware downloaded to firmware slot 6. If no valid firmware revision is present or if this slot is unsupported, all zeros shall be returned this field shall be cleared to 0h.
63:56	Firmware Revision for Slot 7 (FRS7): Contains the revision of the firmware downloaded to firmware slot 7. If no valid firmware revision is present or if this slot is unsupported, all zeros shall be returned this field shall be cleared to 0h.
511:64	Reserved

Modify a portion of Figure 138 (Format NVM) as shown below:

04	Metadata Settings (MSET): This field is set to '1' if the metadata is transferred as part of an extended data LBA. This field is cleared to '0' if the metadata is transferred as part of a separate buffer. The metadata may include protection information, based on the Protection Information (PI) field. If the Metadata Size for the LBA Format selected is 0h, then this field is not applicable.
----	---

Modify a portion of Figure 91 (Identify Power State Descriptor) as shown below:

24	Max Power Scale (MXPS): This field indicates the scale for the Maximum Power field. If this field is cleared to '0', then the scale of the Maximum Power field is in 0.01 Watts. If this field is set to '1', then the scale of the Maximum Power field is in 0.0001 Watts.
----	--

<Note: The entries below apply to NVMe 1.2 and earlier. Similar changes were added as requirements in NVMe 1.2 TP 005>

Modify a portion of 5.10.1.1 (Error Information log) as shown below:

This log page is used to describe extended error information for a command that completed with error or report an error that is not specific to a particular command. Extended error information is provided when the More (M) bit is set to '1' in the Status Field for the completion queue entry associated with the command that completed with error or as part of an asynchronous event with an Error status type. This log page is global to the controller.

This error log may return the last *n* errors. If host software specifies a data transfer of the size of *n* error logs, then the error logs for the **most recent last *n* errors** **is are** returned. The ordering of the entries is based on the time when the error occurred, with the most recent error being returned as the first log **entry**.

Each entry in the log page returned is defined in Figure 78. The log page is a set of 64 byte entries; the **maximum** number of entries supported is indicated in the Identify Controller data structure in Figure 90. **If the log page is full when a new entry is generated, the controller should insert the new entry into the log and discard the oldest entry.**

The controller should clear this log page by removing all entries on power cycle and reset.

Modify a portion of Figure 90 (Identify Controller) as shown below:

262	M	Error Log Page Entries (ELPE): This field indicates the maximum number of Error Information log entries that are stored by the controller. This field is a 0's based value.
-----	---	---

Modify a portion of Figure 78 (Error Information Log) as shown below:

07:00	Error Count: This is a 64-bit incrementing error count, indicating a unique identifier for this error. The error count starts at 1h, is incremented for each unique error log entry, and is retained across power off conditions. A value of 0h indicates an invalid entry; this value is may-be used when there are lost entries or when there are fewer errors than the maximum number of entries the controller supports.
-------	--