



LEGAL NOTICE:

© Copyright 2007 - 2017 NVM Express, Inc. ALL RIGHTS RESERVED.

This erratum to the 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 erratum to the 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 - 2016 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 | 001 |
| Revision Date | 05/15/2017 |
| Affected Spec Ver. | NVM Express™ 1.3 |
| Corrected Spec Ver. | |

Errata Author(s)

| Name | Company |
|--------------|----------|
| Judy Brock | Samsung |
| Mike Allison | SK Hynix |
| Fred Knight | NetApp |
| David Black | Dell EMC |
| Tom Friend | Toshiba |

Errata Overview

Clarified PRP Offset Invalid status code and the PRP Offset field includes bits 1:0.

Clarified Protection Information for Compare when the check is done.

Clarified the fused command error when there is a missing first fused command.

Updated NQN construction example to show 223 bytes as the maximum.

Many clarifications to the Reservation Notification log page, including an incompatible change to address the wrap around condition.

Clarified that Commit Action 03h of the Firmware Commit command may replace existing slot contents before doing activation without reset.

Overlapping Range status usage was clarified.

Revision History

| Revision Date | Change Description |
|---------------|--|
| 04/05/2017 | Initial draft: Clarify PRP Offset Invalid status, Clarify PRP Offset field definition includes bits 1:0 |
| 4/13/2017 | Integration of a few other errata items from the reflector. |
| 4/27/2017 | Refinements from 4/13 meeting. |
| 5/3/2017 | Removed a few items to ECN 002 and finished final revisions of items in ECN 001. |
| 5/11/2017 | Updates based on 5/4 meeting. |
| 5/15/2017 | Minor changes based on 5/11 meeting. |
| | |
| | |
| | |
| | |
| | |

Incompatible Changes

The change to the status code “Command Aborted due to Missing Fused Command” indicates this status code is returned if the previous fused command is missing. This is new functionality not required in NVMe revision 1.3.

The behavior of the Reservation Notification Log has changed for the condition of the maximum number of log pages has been reached. Additionally, the definition of the zero Log Page Count has been modified, where the 0h value now means that the log page is empty (where there were two conflicting interpretations before).

The Controller ID for the Registered Controller data structure in Figure 233 has a new requirement to be FFFFh if it is a dynamic controller that is not associated with a host.

Description of Specification Changes

Modify a portion of Figure 31 (Status Code – Generic Command Status Values) as shown below:

Figure 31: Status Code – Generic Command Status Values

| Value | Description |
|-------|--|
| ... | ... |
| 13h | PRP Offset Invalid: The Offset field for a PRP entry is invalid. This may occur when there is a PRP entry with a non-zero offset after the first entry or when the Offset field in any PRP entry is not Dword aligned (i.e., bits 1:0 are not cleared to 00b). |
| ... | ... |

Modify a portion of Figure 14 (PRP Entry – Page Base Address and Offset) as shown below:


Figure 14: PRP Entry – Page Base Address and Offset

| Bit | Description |
|---------------------------|---|
| 63:02 63:00 | <p>Page Base Address and Offset (PBAO): This field indicates the 64-bit physical memory page address. The lower bits (#:2 n:0) of this field indicate the offset within the memory page. If the memory page size is 4KB, then bits 44:02 11:00 form the Offset; if the memory page size is 8KB, then bits 42:02 12:00 form the Offset, etc. If this entry is not the first PRP entry in the command or a PRP List pointer in a command, then the Offset portion of this field shall be cleared to 0h. The Offset shall be Dword aligned, indicated by bits 1:0 being cleared to 00b.</p> <p>NOTE: The controller shall operate as if bits 1:0 are cleared to 00b. However, the controller is not required to check that bits 1:0 are cleared to 00b. The controller may report an error of PRP Offset Invalid if bits 1:0 are not cleared to 00b.</p> |
| 01:00 | Reserved |

Modify Figure 13 (below) to show the ‘Offset’ field extending through bits 1:0, removing the ‘0 | 0’ fields in the right portion of the figure.

BEFORE:

Figure 13: PRP Entry Layout



AFTER:

Figure 13: PRP Entry Layout



Modify a portion of Figure 200 (Compare – Command Dword 12) as shown below:

Figure 200: Compare – Command Dword 12

| Bit | Description |
|-------|---|
| 31 | Limited Retry (LR): If set to ‘1’, the controller should apply limited retry efforts. If cleared to ‘0’, the controller should apply all available error recovery means to retrieve the data for comparison. |
| 30 | Force Unit Access (FUA): This field specifies that the data read shall be read from non-volatile media. |
| 29:26 | Protection Information Field (PRINFO): Specifies the protection information action and check field, as defined in Error! Reference source not found. The Protection Information Action (PRACT) field shall be cleared to ‘0’. If the Protection Information Check (PRCHK) field is non-zero, a check is performed on the logical block read from NVM (refer to section 8.3.1.4). |
| 25:16 | Reserved |
| 15:00 | Number of Logical Blocks (NLB): This field specifies the number of logical blocks to be compared. This is a 0’s based value. |

Modify a portion of Figure 31 as shown below:

| | |
|-----|---|
| 09h | Command Aborted due to Failed Fused Command: The command was aborted due to the other command in a fused operation failing. |
| 0Ah | Command Aborted due to Missing Fused Command: The command was aborted due to the companion fused command not being found as the subsequent Submission Queue entry. The Submission Queue does not contain the first command followed by the second command for a Fused Operation (refer to Figure 10). |
| 0Bh | Invalid Namespace or Format: The namespace or the format of that namespace is invalid. |

Modify Figure 254 (NQN Construction for Older NVM Subsystems) as shown below:

Figure 254: NQN Construction for Older NVM Subsystems

| Bytes | Description |
|--|---|
| 26:00 | NQN Starting String (NSS): Contains the 27 letter ASCII string "nqn.2014.08.org.nvmexpress:". |
| 30:27 | PCI Vendor ID (VID): Contains the company vendor identifier that is assigned by the PCI SIG as a hexadecimal ASCII string. |
| 34:31 | PCI Subsystem Vendor ID (SSVID): Contains the company vendor identifier that is assigned by the PCI SIG for the subsystem as a hexadecimal ASCII string. |
| 54:35 | Serial Number (SN): Contains the serial number for the NVM subsystem that is assigned by the vendor as an ASCII string. |
| 94:55 | Model Number (MN): Contains the model number for the NVM subsystem that is assigned by the vendor as an ASCII string. |
| 255:95 222:95 | Padding (PAD): Contains spaces (ASCII character 20h). |

Modify section 5.14.1.9.1 (Reservation Notification) as shown below:

The Reservation Notification log page reports one log page from a time ordered queue of reservation notification log pages, if available. A new Reservation Notification log page is created and added to the end of the queue of reservation notifications whenever an unmasked reservation notification occurs on any namespace that is attached to the controller. The Get Log Page command:

- returns a data buffer containing a log page corresponding to the oldest log page in the a-single reservation notification queue (i.e., the log page containing the lowest Log Page Count field; accounting for wrapping); and
- removes that Reservation Notification log page from the queue.

If there are no available Reservation Notification log page entries when a Get Log command is issued, then an empty log page (i.e., all fields in the log page set to zero) shall be returned.

If the controller is unable to store a reservation notification in the Reservation Notification log due to the size of the queue, that reservation notification is lost. If a reservation notification is lost, then the controller shall increment the Log Page Count field of the last reservation notification in the queue (i.e., the Log Page Count field in the last reservation notification in the queue shall contain the value associated with the most recent reservation notification that has been lost).

< ADD BLANK LINE >

The format of the log page is defined in Figure 103. This log page is global to the controller.

Figure 103: Get Log Page – Reservation Notification Log

| Bytes | Description | | | | | | | | | | | | |
|-------|---|-------|------------|---|--|---|------------------------|---|----------------------|---|-----------------------|-------|----------|
| 07:00 | <p>Log Page Count: This is a 64-bit incrementing Reservation Notification log page count, indicating a unique identifier (modulo 64-bit) for this notification. The count starts at 0h following a controller reset, is incremented with each unique log entry for every event that causes a reservation notification regardless of whether that notification is added to the queue, and rolls over to zero one when the maximum count is reached and a new log page is created. If there are no Reservation Notification log pages to return (i.e., the queue of reservation log pages is empty), then this field shall return the value 0h. Subsequent reservation notifications continue incrementing this unique identifier from the last non-zero value (i.e., the value that identified the previous Reservation Notification log page). A value of 0h indicates an empty log entry the log page is empty.</p> | | | | | | | | | | | | |
| 08 | <p>Reservation Notification Log Page Type: This field indicates the Reservation Notification type described by this log page.</p> <table border="1"> <thead> <tr> <th>Value</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Empty Log Page: Get Log Page command was processed when no unread Reservation Notification log pages were available. All the fields of an empty log page shall have a value of zero.</td> </tr> <tr> <td>1</td> <td>Registration Preempted</td> </tr> <tr> <td>2</td> <td>Reservation Released</td> </tr> <tr> <td>3</td> <td>Reservation Preempted</td> </tr> <tr> <td>255:4</td> <td>Reserved</td> </tr> </tbody> </table> | Value | Definition | 0 | Empty Log Page: Get Log Page command was processed when no unread Reservation Notification log pages were available. All the fields of an empty log page shall have a value of zero. | 1 | Registration Preempted | 2 | Reservation Released | 3 | Reservation Preempted | 255:4 | Reserved |
| Value | Definition | | | | | | | | | | | | |
| 0 | Empty Log Page: Get Log Page command was processed when no unread Reservation Notification log pages were available. All the fields of an empty log page shall have a value of zero. | | | | | | | | | | | | |
| 1 | Registration Preempted | | | | | | | | | | | | |
| 2 | Reservation Released | | | | | | | | | | | | |
| 3 | Reservation Preempted | | | | | | | | | | | | |
| 255:4 | Reserved | | | | | | | | | | | | |
| 09 | <p>Number of Available Log Pages: This field indicates the number of additional available Reservation Notification log pages (i.e., the number of unread log pages not counting this one). If there are more than 255 additional available log pages, then a value of 255 is returned. A value of zero indicates that there are no additional available log pages.</p> | | | | | | | | | | | | |
| 11:10 | Reserved | | | | | | | | | | | | |
| 15:12 | <p>Namespace ID: This field indicates the namespace ID of the namespace associated with the Reservation Notification described by this log page.</p> | | | | | | | | | | | | |
| 63:16 | Reserved | | | | | | | | | | | | |

Modify Figure 233 (Registered Controller Data Structure) as shown below:

Figure 233: Registered Controller Data Structure

| Bytes | Description |
|-------|--|
| 1:0 | <p>Controller ID (CNTLID): This field contains the controller ID (i.e., the value of the CNTLID field in the Identify Controller data structure) of the controller whose status is reported in this data structure.</p> <p>If the controller is a dynamic controller (refer to the NVMe over Fabrics specification) that is not associated with a host, then the Controller ID field shall be set to FFFFh.</p> |
| 2 | <p>Reservation Status (RCSTS): This field indicates the reservation status of the controller described by this data structure.</p> <p>Bits 7:1 are reserved</p> <p>Bit 0 is set to '1' if the controller is associated with a host that holds a reservation on the namespace.</p> |
| 7:3 | Reserved |
| 15:8 | <p>Host Identifier (HOSTID): This field contains the 64-bit Host Identifier of the controller described by this data structure.</p> |
| 23:16 | <p>Reservation Key (RKEY): This field contains the reservation key of the host associated with the controller described by this data structure.</p> |

Modify Figure 76 (Firmware Commit – Command Dword 10) as shown below

Figure 76: Firmware Commit – Command Dword 10

| Bit | Description | | | | | | | | | | | | | | | | |
|----------|--|-------|------------|------|--|------|---|------|--|------|--|----------|----------|------|--|------|--|
| 31 | Boot Partition ID (BPID): Specifies the Boot Partition that shall be used for the Commit Action, if applicable. | | | | | | | | | | | | | | | | |
| 30:06 | Reserved | | | | | | | | | | | | | | | | |
| 05:03 | <p>Commit Action (CA): This field specifies the action that is taken on the image downloaded with the Firmware Image Download command or on a previously downloaded and placed image. The actions are indicated in the following table.</p> <table border="1"> <thead> <tr> <th>Value</th> <th>Definition</th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>Downloaded image replaces the existing image, if any, in the specified by the Firmware Slot field. The newly placed is image is not activated.</td> </tr> <tr> <td>001b</td> <td>Downloaded image replaces the existing image, if any, in the specified by the Firmware Slot field. The newly placed is image is activated at the next reset.</td> </tr> <tr> <td>010b</td> <td>The existing image in the specified by the Firmware Slot field is activated at the next reset.</td> </tr> <tr> <td>011b</td> <td>Downloaded image replaces the existing image, if any, in the specified Firmware Slot and is then activated immediately. If there is not a newly downloaded image, then the existing image in the specified firmware slot is activated immediately. The image specified by the Firmware Slot field is requested to be activated immediately without reset.</td> </tr> <tr> <td>100-101b</td> <td>Reserved</td> </tr> <tr> <td>110b</td> <td>Downloaded image replaces the Boot Partition specified by the Boot Partition ID field.</td> </tr> <tr> <td>111b</td> <td>Mark the Boot Partition specified in the BPID field as active and update BPINFO.ABPID.</td> </tr> </tbody> </table> | Value | Definition | 000b | Downloaded image replaces the existing image, if any, in the specified by the Firmware Slot field. The newly placed is image is not activated. | 001b | Downloaded image replaces the existing image, if any, in the specified by the Firmware Slot field. The newly placed is image is activated at the next reset. | 010b | The existing image in the specified by the Firmware Slot field is activated at the next reset. | 011b | Downloaded image replaces the existing image, if any, in the specified Firmware Slot and is then activated immediately. If there is not a newly downloaded image, then the existing image in the specified firmware slot is activated immediately. The image specified by the Firmware Slot field is requested to be activated immediately without reset. | 100-101b | Reserved | 110b | Downloaded image replaces the Boot Partition specified by the Boot Partition ID field. | 111b | Mark the Boot Partition specified in the BPID field as active and update BPINFO.ABPID. |
| Value | Definition | | | | | | | | | | | | | | | | |
| 000b | Downloaded image replaces the existing image, if any, in the specified by the Firmware Slot field. The newly placed is image is not activated. | | | | | | | | | | | | | | | | |
| 001b | Downloaded image replaces the existing image, if any, in the specified by the Firmware Slot field. The newly placed is image is activated at the next reset. | | | | | | | | | | | | | | | | |
| 010b | The existing image in the specified by the Firmware Slot field is activated at the next reset. | | | | | | | | | | | | | | | | |
| 011b | Downloaded image replaces the existing image, if any, in the specified Firmware Slot and is then activated immediately. If there is not a newly downloaded image, then the existing image in the specified firmware slot is activated immediately. The image specified by the Firmware Slot field is requested to be activated immediately without reset. | | | | | | | | | | | | | | | | |
| 100-101b | Reserved | | | | | | | | | | | | | | | | |
| 110b | Downloaded image replaces the Boot Partition specified by the Boot Partition ID field. | | | | | | | | | | | | | | | | |
| 111b | Mark the Boot Partition specified in the BPID field as active and update BPINFO.ABPID. | | | | | | | | | | | | | | | | |
| 02:00 | Firmware Slot (FS): Specifies the firmware slot that shall be used for the Commit Action, if applicable. If the value specified is 0h, then the controller shall choose the firmware slot (slot 1 – 7) to use for the operation. | | | | | | | | | | | | | | | | |

Modify a portion of section 5.12 (Firmware Image Download command) as shown below:

5.12 Firmware Image Download command

...

The host software ~~shall~~ **should** ensure that image pieces do not have Dword ranges that overlap. Firmware portions may be submitted out of order to the controller. Host software shall submit image portions in order when updating a Boot Partition. ~~If ranges overlap, the controller may return an error of Overlapping Range.~~

Modify a portion of Figure 176 (Format NVM – Command Dword 10) as shown below:

| | |
|----|---|
| 08 | Protection Information Location (PIL): If set to '1' and protection information is enabled, then protection information is transferred as the first eight bytes of metadata. If cleared to '0' and protection information is enabled, then protection information is transferred as the last eight bytes of metadata. This setting is reported in the End-to-end Data Protection Type Settings (DPS) Formatted LBA Size field of the Identify Namespace data structure and is constrained by the End-to-end Data Protection Capabilities (DPC) field of the Identify Namespace data structure. |
|----|---|

