

## NVM Express™ Technical Errata

<b>Errata ID</b>	<b>001</b>
<b>Revision Date</b>	<b>1/16/2015</b>
<b>Affected Spec Ver.</b>	<b>NVM Express™ 1.2</b>
<b>Corrected Spec Ver.</b>	

### Errata Author(s)

Name	Company
John Carroll	Intel
Jason Gao	HGST
Peter Onufryk	PMC
Dennis O'Connor	Micron
Dave Landsman	SanDisk
Koichi Nagai	Toshiba
Fred Au	Toshiba
Jim Hatfield	Seagate

### Errata Overview

Clarified protection information pictures and text, corrected various typos, and added clarification on SGL types.

## Revision History

Revision Date	Change Description
9/11/2014	First draft.
9/25/2014	Removed items included in NVMe 1.2 release candidate; added clarification on deleting all namespaces
11/12/2014	Added additional feedback on SGL entries, cleaned up CDW0 fields and opcodes.
11/20/2014	Added changes based on 11/13 workgroup feedback. Approved in 11/20 meeting for 30 day review.
1/16/2015	Ratified.

## Description of Specification Changes

**Modify Figure 124 as shown below:**

**Figure 124: Autonomous Power State Transition – Data Structure Entry**

Bit	Description
63:32	Reserved
31:08	<b>Idle Time Prior to Transition (ITPT):</b> This field specifies the amount of idle time that occurs in this power state prior to transitioning to the Idle Transition Power State. The time is specified in milliseconds. A value of 0h disables the autonomous power state transition feature for this power state.
07:03	<b>Idle Transition Power State (ITPS):</b> This field specifies the power state for the controller to autonomously transition to after there is a continuous period of idle time in the current power state that exceeds time specified in the Idle Time Prior to Transition field. The field specified is required to be a non-operational state as described in Figure 91. <b>This field should not specify a power state with higher reported idle power than the current power state.</b>
02:00	Reserved

**Modify the third paragraph section 7.8 as shown below:**

The default value for each Feature is vendor specific and **set by the manufacturer unless otherwise specified; it** is not changeable. The saveable value is the value that the Feature has after a power on or reset event. The controller may not support a saveable value for a Feature; this is discovered by using the 'supported capabilities' value in the Select field in Get Features. If the controller does not support a saveable value for a Feature, then the default value is used after a power on or reset event. The current value is the value actively in use by the controller for a Feature after a Set Features command completes.

**Modify a portion of section 5.14.1 as shown below:**

#### 5.14.1 Feature Specific Information

Figure 108 defines the Features that may be configured with Set Features and retrieved with Get Features. Figure 109 defines Features that are specific to the NVM Command Set. Some Features utilize a memory buffer to configure or return attributes for a Feature, whereas others only utilize a Dword in the command or completion queue entry. Feature values that are not persistent across power states are reset to their default values as part of a controller reset operation. The default value for each Feature is vendor specific and **set by the manufacturer unless otherwise specified**; it is not changeable. For more information on Features, including default, saveable, and current value definitions, refer to section 7.8.

**Modify a portion of section 3.1.5 as shown below:**

00	RW	0	<p><b>Enable (EN):</b> 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 <b>5.14.1</b> 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 7.3 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 section 5.13 as shown below:**

The Namespace Identifier (CDW1.NSID) field is used as follows for create and delete operations:

- Create: The CDW1.NSID field is reserved for this operation; host software shall set this field to a value of 0h. The controller shall select the next available Namespace Identifier to use for the operation.
- Delete: This field specifies the previously created namespace to delete in this operation. **Specifying a value of FFFFFFFFh may be used to delete all namespaces accessible from a controller.**

**Modify a portion of section 4.9 as shown below:**

A Controller List, defined in Figure 37, is an ordered list of **ascending** controller IDs. The controller identifier is defined in bytes 79:78 of the Identify data structure in Figure 90. Unused entries are zero filled.

**Modify a portion of section 5.13 as shown below:**

The Namespace Identifier (CDW1.NSID) field is used as follows for create and delete operations:


- Create: The CDW1.NSID field is reserved for this operation; host software shall set this field to a value of 0h. The controller shall select ~~the next~~ an available Namespace Identifier to use for the operation.

**Modify Figure 10 as shown below:**


**Figure 10: Command Dword 0**

Bit	Description										
31:16	<b>Command Identifier (CID):</b> This field specifies a unique identifier for the command when combined with the Submission Queue identifier.										
15:14	<b>PRP or SGL for Data Transfer (PSDT):</b> This field specifies whether PRPs or SGLs are used for any data transfer associated with the command. PRPs shall be used for all Admin commands. The definition is described in the table below. <table><tr><th>Value</th><th>Definition</th></tr><tr><td>00b</td><td>PRPs are used for this transfer.</td></tr><tr><td>01b</td><td>SGLs are used for this transfer. Metadata Pointer (MPTR) contains an address of a single contiguous physical buffer that is byte aligned.</td></tr><tr><td>10b</td><td>SGLs are used for this transfer. Metadata Pointer (MPTR) contains an address of an SGL segment containing exactly one SGL Descriptor that is Qword aligned.</td></tr><tr><td>11b</td><td>Reserved</td></tr></table> <p>If there is metadata that is not interleaved with the logical block data, as specified in the Format NVM command, then the Metadata Pointer (MPTR) field is used to point to the metadata. The definition of the Metadata Pointer field is dependent on the setting in this field. Refer to Figure 12.</p>	Value	Definition	00b	PRPs are used for this transfer.	01b	SGLs are used for this transfer. Metadata Pointer (MPTR) contains an address of a single contiguous physical buffer that is byte aligned.	10b	SGLs are used for this transfer. Metadata Pointer (MPTR) contains an address of an SGL segment containing exactly one SGL Descriptor that is Qword aligned.	11b	Reserved
Value	Definition										
00b	PRPs are used for this transfer.										
01b	SGLs are used for this transfer. Metadata Pointer (MPTR) contains an address of a single contiguous physical buffer that is byte aligned.										
10b	SGLs are used for this transfer. Metadata Pointer (MPTR) contains an address of an SGL segment containing exactly one SGL Descriptor that is Qword aligned.										
11b	Reserved										
13:10	Reserved										
09:08	<b>Fused Operation (FUSE):</b> In a fused operation, a complex command is created by “fusing” together two simpler commands. Refer to section 6.1. This field specifies whether this command is part of a fused operation and if so, which command it is in the sequence. <table><tr><th>Value</th><th>Definition</th></tr><tr><td>00b</td><td>Normal operation</td></tr><tr><td>01b</td><td>Fused operation, first command</td></tr><tr><td>10b</td><td>Fused operation, second command</td></tr><tr><td>11b</td><td>Reserved</td></tr></table>	Value	Definition	00b	Normal operation	01b	Fused operation, first command	10b	Fused operation, second command	11b	Reserved
Value	Definition										
00b	Normal operation										
01b	Fused operation, first command										
10b	Fused operation, second command										
11b	Reserved										
07:00	<b>Opcode (OPC):</b> This field specifies the opcode of the command to be executed.										


Modify Figure 213 as shown below (for inclusion in specification – complete update of figures in Visio):




(a) No Protection Information



~~(a)~~ No Protection Information with metadata




(b) Protection Information with PRACT bit cleared to '0' (i.e., pass)




(c) Protection Information with PRACT bit set to '1' (i.e., insert)


Modify Figure 214 as shown below :




(a) No Protection Information



(b) No Protection Information with metadata



(b) Protection Information with PRACT bit cleared to '0' (i.e., pass)



(c) Protection Information with PRACT bit set to '1' (i.e., strip)

Modify a portion of Figure 142 as shown below:

Figure 1: Security Receive – Command Dword 10

Bit	Description
31:24	<b>Security Protocol (SECP):</b> This field specifies the security protocol as defined in SPC-4. The controller shall fail the command with Invalid Parameter indicated if an unsupported value of the Security Protocol is specified.
23:08	<b>SP Specific 1 (SPSP1):</b> The value of this field contains bits 15:08 of the Security Protocol Specific field as defined in SPC-4.
15:08	<b>SP Specific 0 (SPSP0):</b> The value of this field contains bits 07:00 of the Security Protocol Specific field as defined in SPC-4.
07:00	<b>NVMe Security Specific Field (NSSF):</b> Refer to Figure 144 for definition of this field for Security Protocol EAh. For all other Security Protocols this field is reserved.

Modify a portion of Figure 147 as shown below:

**Figure 2: Security Send – Command Dword 10**

Bit	Description
31:24	<b>Security Protocol (SECP):</b> This field specifies the security protocol as defined in SPC-4. The controller shall fail the command with Invalid Parameter indicated if a reserved value of the Security Protocol is specified.
23:08 <del>16</del>	<b>SP Specific 1 (SPSP1):</b> The value of this field <del>contains bits 15:08 of the Security Protocol</del> <b>Specific field is specific to the Security Protocol</b> as defined in SPC-4.
15:08	<b>SP Specific 0 (SPSP0):</b> The value of this field <del>contains bits 07:00 of the Security Protocol</del> <b>Specific field as defined in SPC-4.</b>
07:00	<b>NVMe Security Specific Field (NSSF):</b> Refer to Figure 144 for definition of this field for Security Protocol EAh. For all other Security Protocols this field is reserved.

**Modify a portion of 5.7.1 as shown below:**

A completion queue entry is posted to the Admin Completion Queue if the controller has completed the requested action (specified in the Commit Action field). For requests that specify activation of a new firmware image and return with status code value of 00h, any Controller Level Reset defined in section 7.3.2 activates the specified firmware. **The controller may complete this command with the generic command status value for Internal Error if there are errors storing the firmware image.** Firmware Commit command specific status values are defined in Figure 63.

**Modify a portion of Figure 40 as shown below:**

**Figure 40: Opcodes for Admin Commands**

Opcode by Field			Combined Opcode <sup>2</sup>	O/M <sup>1</sup>	Namespace Identifier Used <sup>3</sup>	Command
Opcode (07)	Opcode (06:02)	Opcode (01:00)				
Generic Command	Function	Data Transfer <sup>4</sup>				
0b	000 00b	00b	00h	M	No	Delete I/O Submission Queue
0b	000 00b	01b	01h	M	No	Create I/O Submission Queue
0b	000 00b	10b	02h	M	Yes	Get Log Page
0b	000 01b	00b	04h	M	No	Delete I/O Completion Queue
0b	000 01b	01b	05h	M	No	Create I/O Completion Queue
0b	000 01b	10b	06h	M	Yes	Identify
0b	000 10b	00b	08h	M	No	Abort
0b	000 10b	01b	09h	M	Yes	Set Features
0b	000 10b	10b	0Ah	M	Yes	Get Features
0b	000 11b	00b	0Ch	M	No	Asynchronous Event Request
0b	000 11b	01b	0Dh	O	Yes	Namespace Management
0b	001 00b	00b	10h	O	No	Firmware Commit
0b	001 00b	01b	11h	O	No	Firmware Image Download
0b	001 01b	01b	15h	O	Yes	Namespace Attachment
<b>I/O Command Set Specific</b>						
1b	na	Na	80h – BFh	O		I/O Command Set specific
<b>Vendor Specific</b>						
1b	na	Na	C0h – FFh	O		Vendor specific

NOTES:

- O/M definition: O = Optional, M = Mandatory.
- Opcodes not listed are reserved.
- A subset of commands uses the Namespace Identifier field (CDW1.NSID). When not used, the field shall be cleared to 0h.
- 00b = no data transfer; 01b = host to controller; 10b = controller to host; 11b = bidirectional

Modify a portion of Figure 149 as shown below:

Figure 3: Opcodes for NVM Commands

Opcode by Field			Combined Opcode <sup>2</sup>	O/M <sup>1</sup>	Command <sup>3</sup>
Opcode (07)	Opcode (06:02)	Opcode (01:00)			
Standard Command	Function	Data Transfer <sup>5</sup>			
0b	000 00b	00b	00h	M	<u>Flush</u>
0b	000 00b	01b	01h	M	<u>Write</u>
0b	000 00b	10b	02h	M	<u>Read</u>
0b	000 01b	00b	04h	O	<u>Write Uncorrectable</u>
0b	000 01b	01b	05h	O	<u>Compare</u>
0b	000 10b	00b	08h	O	Write Zeroes
0b	000 10b	01b	09h	O	<u>Dataset Management</u>
0b	000 11b	01b	0Dh	O <sup>4</sup>	Reservation Register
0b	000 11b	10b	0Eh	O <sup>4</sup>	Reservation Report
0b	001 00b	01b	11h	O <sup>4</sup>	Reservation Acquire
0b	001 01b	01b	15h	O <sup>4</sup>	Reservation Release
<b>Vendor Specific</b>					
1b	na	na	80h – FFh	O	Vendor specific

NOTES:

1. O/M definition: O = Optional, M = Mandatory.
2. Opcodes not listed are reserved.
3. All NVM commands use the Namespace Identifier field (CDW1.NSID).
4. Mandatory if reservations are supported as indicated in the Identify Controller data structure.
5. 00b = no data transfer; 01b = host to controller; 10b = controller to host; 11b = bidirectional

Modify a portion of Figure 79 as shown below:

143:128	<b>Power On Hours:</b> Contains the number of power-on hours. This <del>does</del> may not include time that the controller was powered and in <del>a low power state condition</del> a non-operational power state.
---------	--

Modify a portion of Figure 12 as shown below:

23:16	<p><b>Metadata Pointer (MPTR):</b> This field is valid only if the command has metadata that is not interleaved with the logical block data, as specified in the Format NVM command.</p> <p>If CDW0[15:14] is set to 00b, then this field shall contain the address of a contiguous physical buffer of metadata and shall be Dword aligned.</p> <p>If CDW0[15:14] is set to 01b, then this field shall contain the address of a contiguous physical buffer of metadata and shall be byte aligned.</p> <p>If CDW0[15:14] is set to 10b, then this field shall contain the address of an SGL segment containing exactly one SGL Descriptor and shall be Qword aligned. <b>If the SGL segment is a Data Block descriptor, then it describes the entire data transfer.</b> Refer to section 4.4.</p>
-------	--



Modify a portion of Figure 156 as shown below:

Figure 156: Protection Information Field Definition

Bit	Description
03	<b>Protection Information Action (PRACT):</b> The protection information action field indicates the action to take for the protection information. <del>If this field is set to '1', the protection information is stripped (read) or inserted (write). If this field is cleared to '0', the protection information is passed.</del> This field is only used if the namespace is formatted to use end-to-end protection information. Refer to section 8.3.