



LEGAL NOTICE:

© **Copyright 2008 to 2023 NVM Express®, Inc. ALL RIGHTS RESERVED.**

This 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 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: “© 2008 to 2023 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.

The NVM Express® design mark is a registered trademark of NVM Express, Inc.

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

Technical input submitted to NVM Express® is subject to the terms of the NVM Express® Participant's agreement. Copyright © 2008 to 2023 NVM Express, Inc.

NVM Express® Technical Proposal (TP)

Technical Proposal ID	6032 Out-of-Band Admin Command While Shutdown Enhancement
Revision Date	2023.03.08
Builds on Specification(s)	NVM Express Base Specification 2.0c NVM Express Management Interface 1.2c
References	

Technical Proposal Author(s)

Name	Company
Rory Harris	Intel
Austin Bolen	Dell

Technical Proposal Overview

This proposal:

- Defines the interactions between NVMe Admin Commands processed via the out-of-band mechanism that access media and shutdown.
- Defines a mechanism (Ignore Shutdown) to allow a Management Endpoint to bring media out of a shutdown state to process NVMe Admin Commands that access media via the out-of-band mechanism.
- Defines a new bit (Unsafe Shutdown Reason) in the Persistent Event Log to indicate whether an Unsafe Shutdown was due to a Management Endpoint bringing media out of a shutdown state to process NVMe Admin Commands that access media via the out-of-band mechanism.

Revision History

Revision Date	Change Description
2021.09.23	Initial draft
2021.10.01	Updates after review in Technical Working Group 9-27
2021.10.12	Updates after review in Technical Working Group 10-4
2022.01.20	Added ISH bit to allow OOB to force command executing while controller is shutdown
2022.03.07	Updated shutdown section in NVMe-MI Specification Significantly
2022.03.18	Minor cleanup
2022.04.11	Added names to all added values Clarified language with respect to MMNRRT
2022.04.14	More updates to MMNRRT
2022.04.25	Noted that the media cannot be shutdown using the out-of-band shutdown command after being woken up by out-of-band commands with ISH bit set Added reference to CAP.CRWMS to determine if CRTO.CRWMT is supported
2022.05.01	Removed Change Tracking
2022.05.09	Updated based on Phase 2 exit review feedback. Changed ordering of wording around CAP.CRWMS vs CRTO.CRWMT. Noted that it is outside the scope of the specification whether media is shutdown after processing a new command during shutdown processing
2022.06.30	Updated based on 30-day Member Review Feedback - Formatting changes only
2022.11.01	Removed references to MMNRRT. Will be handled using MPR (Wording is clarified in TP6033)
2022.11.14	Updated to match NVMe Base 2.0c and MI 1.2c
2022.11.15	Fixed field naming in Unsafe Shutdown count definition
2023.01.23	Updated Per 30 day member review feedback.
2023.01.30	Replaced "processed normally" with clarifying language in Table Shutdown Interactions with NVMe Admin Commands that Access Media Updated date copyright to 2023
2023.02.26	Integrated
2023.03.08	Updated per integration review

Description for Changes Document for NVM Express Base Specification 2.0c

Feature Enhancements

- Adds Additional Unsafe Shutdown Logging
 - Adds a new type of Unsafe Shutdown where the controller is in the shutdown state when power is lost but media is not in the shutdown state due to a Management Endpoint processing NVMe Admin Commands that access media with the Ignore Shutdown bit set to '1' via the out-of-band mechanism.
 - Adds a new bit (Unsafe Shutdown Reason) to the Persistent Event Log to indicate whether an Unsafe Shutdown was due to a Management Endpoint bringing media out of a shutdown state to process NVMe Admin Commands that access media with the Ignore Shutdown bit set to '1' via the out-of-band mechanism.

Description for Changes Document for NVM Express Management Interface 1.2c

New Features:

- Adds a new bit to NVMe Admin Commands sent via the out-of-band mechanism
 - **New requirement / incompatible change** in Figure 118 (NVMe Admin Command Request Description)
 - Adds the Ignore Shutdown (ISH) bit to allow the Management Endpoint to bring media out of a shutdown state to process NVMe Admin Commands that access media via the out-of-band mechanism.
- Adds new Shutdown Behavior
 - The Controller may abort out-of-band NVMe Admin Commands sent via the out-of-band mechanism that access media so the controller can handle shutdown.

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 for NVM Express Base Specification 2.0c

5.16.1.3 SMART / Health Information (Log Identifier 02h)

Figure 207: Get Log Page – SMART / Health Information Log

...	...
159:144	<p>Unsafe Shutdowns: Contains the number of unsafe shutdowns. This count shall be incremented when if, and only if, main power is lost when:</p> <ul style="list-style-type: none"> a) the controller does not report it is safe to power down (i.e., the CSTS.SHST field is not set to 10b) prior to loss of main power; or b) media is not in a shutdown state because an Admin command that accesses media as defined by Figure 103 was processed via the out-of-band mechanism with the Ignore Shutdown bit set to '1' (refer to the NVM Express Management Interface Specification) while shutdown processing was occurring or had completed (i.e., while the CSTS.SHST field was set to 01b or 10b).
...	...

5.16.1.14.1.5 NVM Subsystem Hardware Error Event (Event Type 05h)

An NVM Subsystem Hardware Error event shall be recorded in the Persistent Event Log when a supported NVM subsystem hardware error event is detected. Which of the NVM subsystem hardware error events are supported is vendor specific. The NVM subsystem hardware error event shall set the Persistent Event Log Event Format Header:

- Event Type field to 05h; and
- Event Type Revision Field to 042h.

Figure 233: NVM Subsystem Hardware Error Event Codes

Code	Description
08h	<p>Unsafe Shutdown: Indicates that the controller incremented the Unsafe Shutdowns field value in the SMART / Health Information Log.</p> <p>The Additional Hardware Error Information field shall be set to the value from the Unsafe Shutdowns field in the SMART / Health Information log at the time of the event.</p> <p>Refer to Figure XXX for the format of the Additional Hardware Error Information field.</p>

Figure XXX: Additional Hardware Error Information for Unsafe Shutdown Errors

Bytes	Value
15:0	<p>Unsafe Shutdowns: This field shall indicate the value from the Unsafe Shutdowns field in the SMART / Health Information log at the time of the event.</p>

Figure XXX: Additional Hardware Error Information for Unsafe Shutdown Errors

Bytes	Value						
16	<p>Unsafe Shutdown Information: This field contains additional information about the unsafe shutdown.</p>						
	<table border="1"> <thead> <tr> <th data-bbox="451 401 602 443">Bit</th> <th data-bbox="602 401 1344 443">Definition</th> </tr> </thead> <tbody> <tr> <td data-bbox="451 443 602 485">7:1</td> <td data-bbox="602 443 1344 485">Reserved</td> </tr> <tr> <td data-bbox="451 485 602 751">0</td> <td data-bbox="602 485 1344 751"> <p>Unsafe Shutdown Due to Out-Of-Band Activity: If the shutdown was unsafe because media was not in a shutdown state when main power was lost because an Admin command that accesses media as defined by Figure 103 was processed via the out-of-band mechanism with the Ignore Shutdown bit set to '1' (refer to the NVM Express Management Interface Specification) while shutdown processing was occurring or had completed (i.e., while the value of the CSTS.SHST field was set to 01b or 10b), then this bit shall be set to '1'; otherwise, this bit shall be cleared to '0'.</p> </td> </tr> </tbody> </table>	Bit	Definition	7:1	Reserved	0	<p>Unsafe Shutdown Due to Out-Of-Band Activity: If the shutdown was unsafe because media was not in a shutdown state when main power was lost because an Admin command that accesses media as defined by Figure 103 was processed via the out-of-band mechanism with the Ignore Shutdown bit set to '1' (refer to the NVM Express Management Interface Specification) while shutdown processing was occurring or had completed (i.e., while the value of the CSTS.SHST field was set to 01b or 10b), then this bit shall be set to '1'; otherwise, this bit shall be cleared to '0'.</p>
	Bit	Definition					
7:1	Reserved						
0	<p>Unsafe Shutdown Due to Out-Of-Band Activity: If the shutdown was unsafe because media was not in a shutdown state when main power was lost because an Admin command that accesses media as defined by Figure 103 was processed via the out-of-band mechanism with the Ignore Shutdown bit set to '1' (refer to the NVM Express Management Interface Specification) while shutdown processing was occurring or had completed (i.e., while the value of the CSTS.SHST field was set to 01b or 10b), then this bit shall be set to '1'; otherwise, this bit shall be cleared to '0'.</p>						

Description of Specification Changes for NVM Express Management Interface 1.2c

6 NVM Express Admin Command Set

...

NVMe Admin Commands over the out-of-band mechanism may target a Controller that is disabled or held in reset by the host. When this occurs, the NVMe Admin Command is processed normally **except as specified in sections 8.1 and 8.X**.

The Request Message format for NVMe Admin Commands is shown in Figure 117 and is described Figure 118.

Figure 118: NVMe Admin Command Request Description

Bytes	Description										
03:00	NVMe-MI Message Header: Refer to section 3.1.										
04	Opcode (OPC): This field specifies the opcode of the command. Refer to the NVM Express Base Specification.										
05	Command Flags (CFLGS): This field specifies flags for the command.										
	<table border="1"> <thead> <tr> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>7:23</td> <td>Reserved</td> </tr> <tr> <td>2</td> <td>Ignore Shutdown (ISH): The effect of this bit is specified in section 8.X. This bit shall have no effect on the value of the CSTS.SHST field (refer to the NVM Express Base Specification).</td> </tr> <tr> <td>1</td> <td>DOFST Valid (DOFSTV): This bit is not used and shall be ignored by the Management Endpoint for implementations compliant with versions of this specification later than 1.1.</td> </tr> <tr> <td>0</td> <td>DLEN Valid (DLENV): This bit is not used and shall be ignored by the Management Endpoint for implementations compliant with versions of this specification later than 1.1.</td> </tr> </tbody> </table>	Bits	Description	7:23	Reserved	2	Ignore Shutdown (ISH): The effect of this bit is specified in section 8.X . This bit shall have no effect on the value of the CSTS.SHST field (refer to the NVM Express Base Specification).	1	DOFST Valid (DOFSTV): This bit is not used and shall be ignored by the Management Endpoint for implementations compliant with versions of this specification later than 1.1.	0	DLEN Valid (DLENV): This bit is not used and shall be ignored by the Management Endpoint for implementations compliant with versions of this specification later than 1.1.
	Bits	Description									
	7:23	Reserved									
	2	Ignore Shutdown (ISH): The effect of this bit is specified in section 8.X . This bit shall have no effect on the value of the CSTS.SHST field (refer to the NVM Express Base Specification).									
1	DOFST Valid (DOFSTV): This bit is not used and shall be ignored by the Management Endpoint for implementations compliant with versions of this specification later than 1.1.										
0	DLEN Valid (DLENV): This bit is not used and shall be ignored by the Management Endpoint for implementations compliant with versions of this specification later than 1.1.										
...	...										

8.X Shutdown Impacts

If shutdown processing is occurring or complete (i.e., the value of the CSTS.SHST field as defined by the NVM Express Base Specification is set to 01b or 10b) on a Controller, then media may be in a shutdown state. If the media is in a shutdown state, then NVMe Admin Commands received over the out-of-band mechanism that require access to media (refer to Admin Commands Permitted to Return a Status Code of Admin Command Media Not Ready in the NVM Express Base Specification) that are processed while shutdown processing is occurring or complete may be impacted by the Ignore Shutdown bit (ISH) as specified in **Figure YYY**.

If an NVMe Admin Command with the ISH bit set to '1' that accesses media is processed while shutdown processing is occurring or complete (i.e., the value of the CSTS.SHST field is set to 01b or 10b), then the media may transition out of the shutdown state. If the NVM Subsystem loses main power while the media is not in the shutdown state, then the Unsafe Shutdown field in the SMART / Health Information log page is incremented (refer to the NVM Express Base Specification).

Technical input submitted to NVM Express® is subject to the terms of the NVM Express® Participant's agreement. Copyright © 2008 to 2023 NVM Express, Inc.

If the media is transitioned out of the shutdown state using an NVMe Admin Command with the ISH bit set to '1', then the NVM Subsystem is permitted to transition the media back into the shutdown state after the NVMe Admin Command is completed. If the media is not transitioned back into the shutdown state and main power is lost while the media is not in the shutdown state, then the Unsafe Shutdown field is incremented. Whether or not an NVM Subsystem transitions the media back into the shutdown state is outside the scope of this specification.

In all cases where an NVMe Admin Command that accesses media is processed while shutdown processing is occurring or complete (i.e., the value of the CSTS.SHST field is set to 01b or 10b) in **Figure YYY**, if the NVMe Admin Command is aborted with the Status field in Completion Queue Entry Dword 3 set to a value of Commands Aborted due to Power Loss Notification or Admin Command Media Not Ready, then media shall not be transitioned out of the shutdown state.

If shutdown processing is occurring or complete (i.e., the value of the CSTS.SHST field is set to 01b or 10b), then:

- there shall be no impact on access to the FRU Information Device; and
- there shall be no impact on the out-of-band mechanism other than for NVMe Admin Commands that access media as described in this section.

If the Controller is in normal operation (i.e., the value of the CSTS.SHST field is cleared to 00b), then:

- there shall be no impact on access to the FRU Information Device; and
- there shall be no impact on the out-of-band mechanism.

If the NVMe Admin Command does not require access to media, then the ISH bit shall have no effect on the processing of the NVMe Admin Command.

Figure YYY: Shutdown Interactions with NVMe Admin Commands that Access Media

CSTS.SHST Field Value at the Time the Controller Started Processing the NVMe Admin Command	ISH Bit Value	
	ISH Bit Cleared to '0'	ISH Bit Set to '1'
00b: Normal operation (i.e. no shutdown has been requested)	If a shutdown has been requested (e.g., the CC.SHN field as defined by the NVM Express Base Specification is set to 01b for normal shutdown or 10b for abrupt shutdown) while the NVMe Admin Command is being processed, then the NVMe Admin Command may be aborted by the controller. If the NVMe Admin Command is aborted, then the Status field in Completion Queue Entry Dword 3 shall be set to a value of Commands Aborted due to Power Loss Notification.	

Figure YYY: Shutdown Interactions with NVMe Admin Commands that Access Media

CSTS.SHST Field Value at the Time the Controller Started Processing the NVMe Admin Command	ISH Bit Value	
	ISH Bit Cleared to '0'	ISH Bit Set to '1'
01b: Shutdown processing occurring	The NVMe Admin Command may be aborted by the controller. If the NVMe Admin Command is aborted, then the Status field in Completion Queue Entry Dword 3 shall be set to a value of Commands Aborted due to Power Loss Notification.	The NVMe Admin Command shall not be aborted by the controller with the Status field in Completion Queue Entry Dword 3 set to Commands Aborted due to Power Loss Notification or Admin Command Media Not Ready. If media is in a shutdown state when the NVMe Admin Command that requires media access is processed, then the media shall be transitioned out of the shutdown state and then the NVMe Admin Command shall be processed. Since the media is required to be transitioned out of the shutdown state, the NVMe Admin Command processing may take longer than normal. If the processing takes longer than the maximum Request-To-Response time, then a More Processing Required Response is transmitted as specified in section 4.1.2.3 .
10b: Shutdown processing complete	The NVMe Admin Command may be aborted by the controller. If the NVMe Admin Command is aborted, then the Status field in Completion Queue Entry Dword 3 shall be set to a value of Admin Command Media Not Ready.	If the value of the CSTS.SHST field at the time the Controller started processing the NVMe Admin Command was set to 01b (i.e. shutdown processing occurring), then the value of the CSTS.SHST field shall transition to a value of 10b (i.e. shutdown processing complete) within the same amount of time as if no NVMe Admin Command with the ISH bit set to '1' had been processed (e.g., within the amount of time indicated by the RTD3 Entry Latency field as defined by the NVM Express Base Specification if the RTD3 Entry Latency is reported).