



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.

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

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

NVM Express® Technical Proposal

Technical Proposal ID	TP 4100 PMR Controller Reset Enhancement
Revision Date	2023.07.05
Builds on Specification(s)	NVM Express Base Specification 2.0c
References	

Technical Proposal Author(s)

Name	Company
Mike Allison, Kapil Karkra, Bryan Veal	Intel
Mike Allison	Samsung

Technical Proposal Overview

This proposal removes the dependency that PMR is disabled on a Controller Reset.

Revision History

Revision Date	Change Description
2021.05.19	Initial draft
2023.01.20	Update to NVM Express Base Specification 2.0c and 2023.
2023.01.31	Added Admin Queue controller properties.
2023.02.02	Edited on NVM Express Technical WG review.
2023.02.23	Phase 3- no changes but date.
2023.04.10	Updated definition of the Enable bit per feedback from Judy Brock.
2023.04.13	Accepted all changes for integration.
2023.06.30	Integrated
2023.07.05	Converted links to text.

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

Description for Changes Document for NVM Express Base Specification 2.0

New Features/Feature Enhancements/Required Changes:

- PMR Controller Reset Enhancement (mandatory if Persistent Memory Region is supported)
 - Description of change.
 - The Persistent Memory Region (PMR) properties are not modified due to a Controller Reset as the intent of Persistent Memory Region is that operate independently of controller being enabled and disabled.
 - The following PMR properties are not modified by a Controller Reset
 - Persistent Memory Region Capabilities (PMRCAP)
 - Persistent Memory Region Control (PMRCTL)
 - Persistent Memory Region Status (PMRSTS)
 - Persistent Memory Region Elasticity Buffer Size (PMREBS)
 - Persistent Memory Region Sustained Write Throughput (PMRSWTP)
 - **New requirement and incompatible change**
 - All Persistent Memory Region properties are not affected by a Controller Reset
 - References
 - Technical Proposal TP4100

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 the NVM Express Base Specification 2.0c

Modify a portions of section 3 as shown below:

3 NVM Express Architecture

3.1 NVM Controller Architecture

...

3.1.3 Controller Properties

...

3.1.3.5 Offset 14h: CC – Controller Configuration

...

Figure 46: Offset 14h: CC – Controller Configuration

Bits	Type	Reset	Description
...			
00	RW	0b	<p>Enable (EN): When set to '1', then the controller shall process commands. When cleared to '0', then the controller shall not process commands nor post completion queue entries to Completion Queues. When this bit transitions from '1' to '0', the controller is reset (i.e., a Controller Reset). That 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. That reset does not affect transport specific state (e.g. PCI Express registers including MMIO MSI-X registers), nor the Admin Queue properties (AQA, ASQ, or ACQ). Refer to section 3.7.2 for the effects of that reset on all All other controller properties defined in this section and. Internal internal controller state (e.g., Feature values defined in section 5.27.1 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 that Controller Reset. Refer to section 3.6.</p> <p>When this bit 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 bit 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 bit from a '0' to a '1' when CSTS.RDY is a '1' or clearing this bit from a '1' to a '0' when CSTS.RDY is cleared to '0' has undefined results. The Admin Queue properties (AQA, ASQ, and ACQ) are only allowed to be modified when this bit is cleared to '0'.</p> <p>If an NVM Subsystem Shutdown is in progress or is completed (i.e., CSTS.ST is set to '1', and CSTS.SHST is set to 01b or 10b), then writes to this field modify the field value but have no effect. Refer to section 3.6.3 for details.</p>

...

3.1.3.8 Offset 24h: AQA – Admin Queue Attributes

This property defines the attributes for the Admin Submission Queue and Admin Completion Queue. The Queue Identifier for the Admin Submission Queue and Admin Completion Queue is 0h. The Admin Submission Queue's priority is determined by the arbitration mechanism selected, refer to section 3.4.4.

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

The Admin Submission Queue and Admin Completion Queue are required to be in physically contiguous memory.

This property shall not be reset by Controller Reset.

...

3.1.3.9 Offset 28h: ASQ – Admin Submission Queue Base Address

This property defines the base memory address of the Admin Submission Queue.

This property shall not be reset by Controller Reset.

...

3.1.3.10 Offset 30h: ACQ – Admin Completion Queue Base Address

This property defines the base memory address of the Admin Completion Queue.

This property shall not be reset by Controller Reset.

...

3.1.3.22 Offset E00h: PMRCAP – Persistent Memory Region Capabilities

This property indicates capabilities of the Persistent Memory Region. If the controller does not support the Persistent Memory Region feature, then this property shall be cleared to 0h.

This property shall not be reset by Controller Reset.

...

3.1.3.23 Offset E04h: PMRCTL – Persistent Memory Region Control

This optional property controls the operation of the Persistent Memory Region. If the controller does not support the Persistent Memory Region feature, then this property shall be cleared to 0h.

This property shall not be reset by Controller Reset.

...

3.1.3.24 Offset E08h: PMRSTS – Persistent Memory Region Status

This optional property provides the status of the Persistent Memory Region. If the controller does not support the Persistent Memory Region feature, then this property shall be cleared to 0h.

This property shall not be reset by Controller Reset.

...

3.1.3.25 Offset E0Ch: PMREBS – Persistent Memory Region Elasticity Buffer Size

This optional property identifies to the host the size of the PMR elasticity buffer. A value of 0h in this property indicates to the host that no information regarding the presence or size of a PMR elasticity buffer is available.

This property shall not be reset by Controller Reset.

...

3.1.3.26 Offset E10h: PMRSWTP – Persistent Memory Region Sustained Write Throughput

This optional property identifies to the host the maximum PMR sustained write throughput. A value of 0h in this property indicates to the host that no information regarding the PMR sustained write throughput is available.

This property shall not be reset by Controller Reset.

...

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

3.7 Resets

...

3.7.2 Controller Level Reset

...

A Controller Level Reset consists of the following actions:

- The controller stops processing any outstanding Admin or I/O commands;
- All I/O Submission Queues are deleted;
- All I/O Completion Queues are deleted;
- The controller is brought to an Idle state. When this is complete, CSTS.RDY is cleared to '0'; and
- All controller properties defined in section 3.1.3 and internal controller state are reset, with the following exceptions:
 - for controllers using a memory-based transport:
 - ~~○ the Admin Queue properties (AQA, ASQ, or ACQ) are not reset as part of a Controller Reset;~~
 - ~~○ the Controller Memory Buffer Memory Space Control property (CMBMSC) is reset as part of neither a Controller Reset nor a Function Level Reset; and~~
 - ~~○ the Persistent Memory Region Memory Space Control Upper property (PMRMSCU) and the Persistent Memory Region Memory Space Control Lower property (PMRMSCL) are not reset as part of a Controller Reset;~~
 - the following are not reset as part of a Controller Level Reset caused by a Controller Reset:
 - Admin Queue properties (i.e., AQA, ASQ, and ACQ);
 - Persistent Memory Region properties (i.e., PMRCAP, PMRCTL, PMRSTS, PMREBS, PMRSWTP, PMRMSCU, and PMRMSCL); and
 - the Controller Memory Buffer Memory Space Control property (CMBMSC);
 - and
 - the following are not reset as part of a Controller Level Reset caused by a Function Level Reset:
 - the Controller Memory Buffer Memory Space Control property (CMBMSC);
 - and
 - for controllers using a message-based transport:
 - there are no exceptions