



LEGAL NOTICE:

© Copyright 2007 - 2020 NVM Express, Inc. **ALL RIGHTS RESERVED.**

This erratum to the NVM Express revision 1.4 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 1.4 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 - 2020 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 Group
3855 SW 153rd Drive
Beaverton, OR 97003 USA
info@nvmexpress.org

NVM Express™ Technical Errata

Errata ID	006
Revision Date	8/27/2020
Affected Spec Ver.	NVM Express™ 1.4a
Corrected Spec Ver.	

Errata Author(s)

Name	Company
Mike Allison, Bryan Veal	Intel

Errata Overview

TP 4054 Added the PMRMSC register. That register is defined as 64-bit and required to be on a qword boundary. During the integration the register was not assigned qword aligned offset not adhering to the approved TP and was published. To address the non-aligned 64-bit register offset, the technical WG decided to split the register into two 32-bit registers. This ECN defines the changes for that decision.

Revision History

Revision Date	Change Description
5/26/2020	Initial creation
5/26/2020	Swapped the register address to align to the original 64-bit address.
5/26/2020	Fixed caption of figure TBD and replaced other references to PMRMSC.
5/28/2020	Use Upper/Lower instead of 1/2. Editorial changes to the host accesses of the registers to specify it is related to the width defined by NVMe. Clarified that the CBA fields should not be changed while CMSE bit is set to '1' for PMR and CMB.
6/2/2020	Updated section 3 wording (avoid using " 's ").
6/3/2020	Made sure updates to CMBMSC is clear that CBA field cannot be set when CMSE bit is set to '1' and continues to be set to '1' when CBA field is updated (i.e. one transaction to set CBA and set CSME bit to '1'). Made sure updates to PMRMSC are clear that CBA field cannot be set when CMSE bit is set to '1' and continues to be set to '1' when CBA field is updated (i.e. two transactions to set PMRMSC.CBA, PMRMSC.CBA and set PMRMSC .CSME bit to '1').
6/4/2020	Technical WG decided the setting of CBA while CSME cleared to '0' was not required. Removed this change that also removed CMB changes.
7/9/2020	Renamed file for membership review. Removed a double space in text.
8/25/2020	Integrated into NVMe Base Specification
8/27/2020	Accepted all changes and removed all comments.

Incompatible Changes

- The PMCMSC register is split into two 32-bit registers.

Description of Specification Changes

Modify a portion of section 3 as shown below:

3 Controller Registers

Controller registers are located in the MLBAR/MUBAR registers (PCI BAR0 and BAR1) that shall be mapped to a memory space that supports in-order access and variable access widths. For many computer architectures, specifying the memory space as uncacheable produces this behavior. The host shall not issue locked accesses. The host shall access a registers ~~in using the width specified for that their register native width~~ or using aligned 32-bit accesses. Violation of either of these host requirements results in undefined behavior.

Accesses that target any portion of two or more registers are not supported.

All reserved registers and all reserved bits within registers are read-only and return 0h when read. Software shall not rely on 0h being returned.

Modify a portion of Figure 68 in section 3.1 as shown below:

3.1 Register Definition

...

Figure 1: Register Definition

Start	End	Symbol	Description
0h	7h	CAP	Controller Capabilities
8h	Bh	VS	Version
Ch	Fh	INTMS	Interrupt Mask Set
10h	13h	INTMC	Interrupt Mask Clear
14h	17h	CC	Controller Configuration
18h	1Bh	Reserved	Reserved
1Ch	1Fh	CSTS	Controller Status
20h	23h	NSSR	NVM Subsystem Reset (Optional)
24h	27h	AQA	Admin Queue Attributes
28h	2Fh	ASQ	Admin Submission Queue Base Address
30h	37h	ACQ	Admin Completion Queue Base Address
38h	3Bh	CMBLOC	Controller Memory Buffer Location (Optional)
3Ch	3Fh	CMBSZ	Controller Memory Buffer Size (Optional)
40h	43h	BPINFO	Boot Partition Information (Optional)
44h	47h	BPRSEL	Boot Partition Read Select (Optional)
48h	4Fh	BPMBL	Boot Partition Memory Buffer Location (Optional)
50h	57h	CMBMSC	Controller Memory Buffer Memory Space Control (Optional)
58h	5Bh	CMBSTS	Controller Memory Buffer Status (Optional)
5Ch	DFFh	Reserved	Reserved
E00h	E03h	PMRCAP	Persistent Memory Capabilities (Optional)
E04h	E07h	PMRCTL	Persistent Memory Region Control (Optional)
E08h	E0Bh	PMRSTS	Persistent Memory Region Status (Optional)
E0Ch	E0Fh	PMREBS	Persistent Memory Region Elasticity Buffer Size
E10h	E13h	PMRSWTP	Persistent Memory Region Sustained Write Throughput

Figure 1: Register Definition

Start	End	Symbol	Description
E14h	E17h	PMRMSC_L	Persistent Memory Region Controller Memory Space Control Lower (Optional)
E18h	E1Bh	PMRMSC_U	Persistent Memory Region Controller Memory Space Control Upper (Optional)
E1Ch	FFFh	Reserved	Command Set Specific
1000h	1003h	SQ0TDBL	Submission Queue 0 Tail Doorbell (Admin)
1000h + (1 * (4 << CAP.DSTRD))	1003h + (1 * (4 << CAP.DSTRD))	CQ0HDBL	Completion Queue 0 Head Doorbell (Admin)

Modify a portion of section Figure 93 in section 3.1.20 as shown below:

3.1.20 Offset E08h: PMRSTS – Persistent Memory Region Status

...

Figure 93: Offset E08h: PMRSTS – Persistent Memory Region Status

Bits	Type	Reset	Description												
31:13	RO	0h	Reserved												
12	RO	0b	Controller Base Address Invalid (CBAI): This field indicates whether the controller has failed to enable the Persistent Memory Region's controller memory space because the controller 64-bit base address specified by PMRMSCU.CBA and PMRMSC.L.CBA PMRMSC.CBA is invalid. If PMRCAP.CMSS is set to '1', PMRMSC.CMSE PMRMSC.L.CMSE is set to '1', and the controller 64-bit base address specified by PMRMSCU.CBA and PMRMSC.L.CBA PMRMSC.CBA is invalid, this bit shall be set to '1'. Otherwise, this bit shall be cleared to '0'.												
11:9	RO	000b	Health Status (HSTS): If the Persistent Memory Region is ready, then this field indicates the health status of the Persistent Memory Region. This field is always cleared to 000b when the Persistent Memory Region is not ready.												
			The health status values are defined as:												
			<table><tr><th>Value</th><th>Definition</th></tr><tr><td>000b</td><td>Normal Operation: The Persistent Memory Region is operating normally.</td></tr><tr><td>001b</td><td>Restore Error: The Persistent Memory Region is operating normally and is persistent; however, the contents of the Persistent Memory Region may not have been restored correctly (i.e., may not contain the contents prior to the last power cycle, NVM subsystem reset, Controller Level Reset, or Persistent Memory Region disable).</td></tr><tr><td>010b</td><td>Read Only: The Persistent Memory Region is read only. PCI Express memory write requests do not update the Persistent Memory Region. PCI Express memory read requests to the Persistent Memory Region return correct data.</td></tr><tr><td>011b</td><td>Unreliable: The Persistent Memory Region has become unreliable. PCI Express memory reads may return invalid data or generate poisoned PCI Express TLP(s). Persistent Memory Region memory writes may not update memory or may update memory with undefined data. The Persistent Memory Region may also have become non-persistent.</td></tr><tr><td>100b to 111b</td><td>Reserved</td></tr></table>	Value	Definition	000b	Normal Operation: The Persistent Memory Region is operating normally.	001b	Restore Error: The Persistent Memory Region is operating normally and is persistent; however, the contents of the Persistent Memory Region may not have been restored correctly (i.e., may not contain the contents prior to the last power cycle, NVM subsystem reset, Controller Level Reset, or Persistent Memory Region disable).	010b	Read Only: The Persistent Memory Region is read only. PCI Express memory write requests do not update the Persistent Memory Region. PCI Express memory read requests to the Persistent Memory Region return correct data.	011b	Unreliable: The Persistent Memory Region has become unreliable. PCI Express memory reads may return invalid data or generate poisoned PCI Express TLP(s). Persistent Memory Region memory writes may not update memory or may update memory with undefined data. The Persistent Memory Region may also have become non-persistent.	100b to 111b	Reserved
			Value	Definition											
			000b	Normal Operation: The Persistent Memory Region is operating normally.											
			001b	Restore Error: The Persistent Memory Region is operating normally and is persistent; however, the contents of the Persistent Memory Region may not have been restored correctly (i.e., may not contain the contents prior to the last power cycle, NVM subsystem reset, Controller Level Reset, or Persistent Memory Region disable).											
			010b	Read Only: The Persistent Memory Region is read only. PCI Express memory write requests do not update the Persistent Memory Region. PCI Express memory read requests to the Persistent Memory Region return correct data.											
011b	Unreliable: The Persistent Memory Region has become unreliable. PCI Express memory reads may return invalid data or generate poisoned PCI Express TLP(s). Persistent Memory Region memory writes may not update memory or may update memory with undefined data. The Persistent Memory Region may also have become non-persistent.														
100b to 111b	Reserved														

Figure 93: Offset E08h: PMRSTS – Persistent Memory Region Status

Bits	Type	Reset	Description
8	RO	0b	Not Ready (NRDY): This bit indicates if the Persistent Memory Region is ready for use. If this bit is cleared to '0' and the PMRCTL.EN is set to '1', then the Persistent Memory Region is ready to accept and process PCI Express memory read and write requests. If this bit is set to '1' or the PMRCTL.EN bit is cleared to '0', then the Persistent Memory Region is not ready to process PCI Express memory read and write requests.
7:0	RO	0h	Error (ERR): When the Persistent Memory Region is ready and operating normally, this field indicates whether previous memory writes to the Persistent Memory Region have completed without error. If this field is cleared to 0h, then previous writes to the Persistent Memory Region have completed without error and that the values written are persistent. A non-zero value in this field indicates the occurrence of an error that may have caused one or more of the previous writes to not have completed successfully. The meaning of any particular non-zero value is vendor specific. Once this field takes on a non-zero value, it maintains a non-zero value until the PCI Function is reset.

Modify a portion of section 3.1.23 as shown below:

3.1.23 Offset E14h: PMRMSC~~L~~– Persistent Memory Region Memory Space Control **Lower**

This register ~~and the PMRMSCU register specifies specify~~ how the controller references the Persistent Memory Region with host-supplied addresses. If the controller supports the Persistent Memory Region's controller memory space (PMRCAP.CMSS), this register is mandatory. Otherwise, this register is reserved. **The host shall access this register with aligned 32-bit accesses.**

This register shall not be reset by Controller Reset.

Figure 96: Offset E14h: PMRMSC~~L~~ – Persistent Memory Region Memory Space Control **Lower**

Bits	Type	Reset	Description
63 31:12	RW	0h	Controller Base Address (CBA): This field specifies the 20 least significant bits of the 52 most significant bits of the 64-bit base address for the Persistent Memory Region's controller address range. The Persistent Memory Region's controller base address and its size determine its controller address range. The specified 64-bit base address specified by this field and PMRMSCU.CBA when the CMSE bit is set to '1' shall be valid only under the following conditions: <ul style="list-style-type: none"> a) no part of the Persistent Memory Region's controller address range is greater than $2^{64} - 1$; and b) if the Controller Memory Buffer's controller memory space is enabled, then the Persistent Memory Region's controller address range does not overlap the Controller Memory Buffer's controller address range.
11:02	RO	0h	Reserved
01	RW	0b	Controller Memory Space Enable (CMSE): This bit specifies whether addresses supplied by the host are permitted to reference the Persistent Memory Region. If this bit is set to '1' and the controller base address is valid, then the Persistent Memory Region's controller memory space is enabled. Otherwise, the controller memory space is not enabled. If the Persistent Memory Region's controller memory space is enabled, then addresses supplied by the host that fall within the Persistent Memory Region's controller address range shall reference the Persistent Memory Region. If the Persistent Memory Region's controller memory space is not enabled, then no address supplied by the host shall reference the Persistent Memory Region. Instead, such addresses shall reference memory spaces other than the Persistent Memory Region.
00	RO	0b	Reserved

Add section 3.1.TBD as shown below:

3.1.TBD Offset E18h: PMRMSCU – Persistent Memory Region Memory Space Control Upper

This register and the PMRMSCU register specify how the controller references the Persistent Memory Region with host-supplied addresses. If the controller supports the Persistent Memory Region's controller memory space (PMRCAP.CMSS), this register is mandatory. Otherwise, this register is reserved. The host shall access this register with aligned 32-bit accesses.

This register shall not be reset by Controller Reset.

Figure TBD: Offset E18h: PMRMSCU – Persistent Memory Region Memory Space Control Upper

Bits	Type	Reset	Description
31:00	RW	0h	Controller Base Address (CBA): This field specifies the 32 most significant bits of the 52 most significant bits of the 64-bit base address for the Persistent Memory Region's controller address range. The Persistent Memory Region's controller base address and its size determine its controller address range.

Modify a portion of section 4.8 as shown below:

4.8 Persistent Memory Region

...

The controller uses the PMR's controller address range to reference PMR with addresses supplied by the host. The PCI Express address range and the controller address range of the PMR may differ, but both ranges have the same size, and equivalent offsets within each range have a one-to-one correspondence. The host configures the controller address range via the ~~PMRMSC register~~ **PMRMSCU and PMRMSCU registers**.

The host enables the PMR's controller memory space via the ~~PMRMSC.CMSE~~ **PMRMSCU.CMSE** bit. When controller memory space is enabled, if host supplies an address referencing the PMR's controller address range, then the controller directs memory read or write requests for this address to the PMR.

Modify a portion of section 7.3.2 as shown below:

7.3.2 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 registers defined in section 3 and internal controller state are reset, except as follows:
 - the Admin Queue registers (AQA, ASQ, or ACQ) are not reset as part of a Controller Reset;
 - the Controller Memory Buffer Memory Space Control register (CMBMSC) is reset as part of neither a Controller Reset nor a Function Level Reset; and
 - the ~~Persistent Memory Region Memory Space Control register (PMRMSC)~~ **Persistent Memory Region Memory Space Control Upper register (PMRMSCU) and the Persistent Memory Region Memory Space Control Lower register (PMRMSCU)** are not reset as part of a Controller Reset.