



#### **LEGAL NOTICE:**

© Copyright 2007 - 2019 NVM Express, Inc. ALL RIGHTS RESERVED.

This NVM Express revision 1.3 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 NVM Express revision 1.3 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: "© 2007 - 2019 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 153<sup>rd</sup> Drive  
Beaverton, OR 97003 USA  
info@nvmexpress.org

## NVM Express Technical Proposal for New Feature

Technical Proposal ID	4051
Change Date	2019-02-19
Builds on Specification(s)	NVM Express 1.3c
Referenced Ratified Technical Proposal(s)	TP 4000a Persistent Memory Region

### Technical Proposal Author(s)

Name	Company
Mike Allison, Jonathan Hughes, Raymond Robles	Intel
Peter Onufryk	Microsemi
Judy Brock	Samsung

This technical proposal relaxes restrictions on the host usage of the Controller Buffer Memory.

### Revision History

Revision Date	Change Description
2018-07-23	Initial version
2018-07-31	Moved new options to CMBLOC register; expanded initial 3 options to 5 options, renaming/word-smithing of new register bits
2018-08-01	Corrected the CBM queue requirement about a queue solely being contained in the CMB.
2018-08-14	Remove the usage of the word employed. Fixed spelling.
2018-08-24	Fixed figures numbers to match NVMe 1.3c. Added I/O Submission Queues and I/O Completion Queues to be dword aligned if located in the Controller Memory Buffer,
2018-11-06	Fixing the functional change made by changing "may only" to "shall" in the changes for CMBLOC.CDPCILS bit as identified by Judy Brock.
2018-11-07	Updating the definition of CDPCILS based on the change from 2018-11-07.
2018-11-08	Wordsmithing the definition of CDPCILS. Corrected the definition of CDPMLS that was accidentally changed with a cut/paste issue.
2018-12-05	More wordsmithing on text associated with CDPCILS
2019-01-15	30 member editorial changes and account for 2019.
2019-01-16	Editorial changes per member review comment: make PRP List and SGL be entirely in CMB or entirely outside CMB; make data and metadata be entirely in CMB or entirely outside CMB.
2019-01-17	Integration
2019-02-19	Ratified

### Description for NVMe Changes Document

- The restrictions on the Controller Memory Buffer in relation to Submission queues, commands, and separation of data and metadata changed to use supported bits in the Controller Memory Buffer Size register. The Controller Memory Buffer section (refer to section 4.7) is modified to apply the

restrictions only if the newly added supported bits indicate the restrictions are enforced. All changes are backwards compatible.

## Description of Specification Changes

The Controller Memory Buffer Size register added bits to control the Controller Memory Buffer restrictions to the Submission Queue, commands, and separation of data and metadata.

### Markup Conventions:

Black: Unchanged (however, hot links are removed)

~~Red-Strikethrough~~: Deleted

Red: New

Red Highlighted: TBD values, anchors, and links to be inserted.

<Green Bracketed>: Notes to editor

**Modify portions of section figure 3.1.11 (Offset 38h: CMBLOC – Controller Memory Buffer Location) as shown below:**

Bits	Type	Reset	Description
31:12	RO	Impl Spec	<b>Offset (OFST):</b> Indicates the offset of the Controller Memory Buffer in multiples of the Size Unit specified in CMBSZ. This value shall be 4KB aligned.

Bits	Type	Reset	Description
11:0309	RO	0h	Reserved
08	RO	Impl Spec	<b>CMB Queue Dword Alignment (CQDA):</b> If this bit is set to '1', then I/O Submission Queues and I/O Completion Queues contained in the Controller Memory Buffer are dword aligned. If this bit is cleared to '0', then the I/O Submission Queues and I/O Completion Queues contained in the Controller Memory Buffer are aligned as defined by the PRP1 field of a Create I/O Submission Queue command (refer to Figure 55) or a Create I/O Completion Queue command (refer to Figure 51).
07	RO	Impl Spec	<b>CMB Data Metadata Mixed Memory Support (CDMMMS):</b> If this bit is set to '1', then the restriction on data and metadata use of Controller Buffer Memory by a command as defined in section 4.7 is not enforced. If this bit is cleared to '0', then the restriction on data and metadata use of Controller Buffer Memory by a command as defined in section 4.7 is enforced.
06	RO	Impl Spec	<b>CMB Data Pointer and Command Independent Locations Support (CDPCILS):</b> If this bit is set to '1', then the restriction that the PRP Lists and SGLs shall not be located in the Controller Buffer Memory if the command that they are associated with is not located in the Controller Buffer Memory is not enforced (refer to section 4.7). If this bit is cleared to '0', then that restriction is enforced.
05	RO	Impl Spec	<b>CMB Data Pointer Mixed Locations Support (CDPMLS):</b> If this bit is set to '1', then the restriction that for a particular PRP List or SGL associated with a single command, all memory that is associated with that particular PRP List or SGL shall reside in either the Controller Memory Buffer or outside the Controller Memory Buffer, is not enforced (refer to section 4.7). If this bit is cleared to '0', then that restriction is enforced.
04	RO	Impl Spec	<b>CMB Queue Physically Discontiguous Support (CQPDS):</b> If this bit is set to '1', then the restriction that for all queues in the Controller Memory Buffer, the queue shall be physically contiguous, is not enforced (refer to section 4.7). If this bit is cleared to '0', then that restriction is enforced.
03	RO	Impl Spec	<b>CMB Queue Mixed Memory Support (CQMMS):</b> If this bit is set to '1', then for a particular queue placed in the Controller Memory Buffer, the restriction that all memory associated with that queue shall reside in the Controller Memory Buffer is not enforced (refer to section 4.7). If this bit is cleared to '0', then that requirement is enforced.
02:00	RO	Impl Spec	<b>Base Indicator Register (BIR):</b> Indicates the Base Address Register (BAR) that contains the Controller Memory Buffer. For a 64-bit BAR, the BAR for the lower 32-bits of the address is specified. Values 0h, 2h, 3h, 4h, and 5h are valid.

### Modify portions of section 4.7 (Controller Memory Buffer) as shown below:

A controller memory based queue is used in the same manner as a host memory based queue – the difference is the memory address used is located within the controller's own memory rather than in the host memory. The Admin or I/O Queues may be placed in the Controller Memory Buffer. If the CMBLOC.CQMMS bit (refer to Figure 82) is cleared to '0', then for ~~For~~ a particular queue, all memory associated with it shall reside in either the Controller Memory Buffer or ~~host memory~~ outside the Controller Memory Buffer.

If the CMBLOC.CQPDS bit (refer to Figure 82) is cleared to '0', then for ~~For~~ all queues in the Controller Memory Buffer, the queue shall be physically contiguous.

The controller may support PRP Lists and SGLs in the Controller Memory Buffer. If the CMBLOC.CDPMLS bit (refer to Figure 82) is cleared to '0', then for ~~For~~ a particular PRP List or SGL associated with a single command, all memory associated with the PRP List or SGLs shall ~~be reside in~~ either entirely located in the Controller Memory Buffer or ~~host memory~~ entirely located outside the Controller Memory Buffer.

~~The~~ PRP Lists and SGLs associated with ~~for~~ a command may ~~only~~ be placed in the Controller Memory Buffer if ~~the associated that~~ command is present in a Submission Queue in the Controller Memory Buffer. If:

- a) CMBLOC.CDPCILS bit (refer to Figure 82) is cleared to '0'; and
- b) a command is not present in a Submission Queue in the Controller Memory Buffer,

then the PRP Lists and SGLs associated with that command shall not be placed in the Controller Memory Buffer.

The controller may support data and metadata in the Controller Memory Buffer. If the CMBLOC.CDMMMS bit (refer to Figure 82) is cleared to '0', then ~~A~~all data and metadata, if any, associated with a particular command shall be either entirely located in the Controller Memory Buffer or ~~entirely located~~ outside the Controller Memory Buffer.

If the requirements for the Controller Memory Buffer use are violated by the host, the controller shall fail the associated command with Invalid Use of Controller Memory Buffer status.

**Modify portions of figure 51 (Create I/O Completion Queue – PRP Entry 1) as shown below:**

63:00	<p><b>PRP Entry 1 (PRP1):</b> If CDW11.PC is set to '1', then this field specifies a 64-bit base memory address pointer of the Completion Queue that is physically contiguous <del>and</del>. The address pointer is memory page aligned (based on the value in CC.MPS) <del>unless otherwise specified</del>. If CDW11.PC is cleared to '0', then this field specifies a PRP List pointer that describes the list of pages that constitute the Completion Queue <del>and</del>. The list of pages is memory page aligned (based on the value in CC.MPS) <del>unless otherwise specified</del>. In both cases the PRP Entry shall have an offset of 0h. In a non-contiguous Completion Queue, each PRP Entry in the PRP List shall have an offset of 0h. If there is a PRP Entry with a non-zero offset, then the controller should return an error of PRP Offset Invalid.</p>
-------	--

**Modify a portion of figure 53 (Create I/O Completion Queue – Command Dword 11) as shown below:**

00	<p><b>Physically Contiguous (PC):</b> If set to '1', then the Completion Queue is physically contiguous and PRP Entry 1 (PRP1) is the address of a contiguous physical buffer. If cleared to '0', then the Completion Queue is not physically contiguous and PRP Entry 1 (PRP1) is a PRP List pointer.</p> <p>If the:</p> <ul style="list-style-type: none"> <li>• queue is located in the Controller Memory Buffer;</li> <li>• <del>and</del> PC is cleared to '0'; and</li> <li>• CMBLOC.CQPDS is cleared to '0',</li> </ul> <p><del>then</del> the controller shall fail the command with Invalid Use of Controller Memory Buffer status.</p>
----	--

**Modify portions of figure 55 (Create I/O Submission Queue – PRP Entry 1) as shown below:**

63:00	<p><b>PRP Entry 1 (PRP1):</b> If CDW11.PC is set to '1', then this field specifies a 64-bit base memory address pointer of the Submission Queue that is physically contiguous <del>and</del>. The address pointer is memory page aligned (based on the value in CC.MPS) <del>unless otherwise specified</del>. If CDW11.PC is cleared to '0', then this field specifies a PRP List pointer that describes the list of pages that constitute the Submission Queue <del>and</del>. The list of pages is memory page aligned (based on the value in CC.MPS) <del>unless otherwise specified</del>. In both cases, the PRP Entry shall have an offset of 0h. In a non-contiguous Submission Queue, each PRP Entry in the PRP List shall have an offset of 0h. If there is a PRP Entry with a non-zero offset, then the controller should return an error of PRP Offset Invalid.</p>
-------	---

**Modify a portion of figure 57 (Create I/O Submission Queue – Command Dword 11) as shown below:**

00	<p><b>Physically Contiguous (PC):</b> If set to '1', then the Submission Queue is physically contiguous and PRP Entry 1 (PRP1) is the address of a contiguous physical buffer. If cleared to '0', then the Submission Queue is not physically contiguous and PRP Entry 1 (PRP1) is a PRP List pointer. If this bit is cleared to '0' and CAP.CQR is set to '1', the controller should return an error of Invalid Field in Command.</p> <p>If the:</p> <ul style="list-style-type: none"> <li>• queue is located in the Controller Memory Buffer;</li> <li>• <del>and</del> PC is cleared to '0'; <del>and</del></li> <li>• CMBLOC.CQPDS is cleared to '0',</li> </ul> <p><del>then</del> the controller shall fail the command with Invalid Use of Controller Memory Buffer status.</p>
----	---