



**LEGAL NOTICE:**

© Copyright 2007 - 2018 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 - 2018 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 Proposal for New Feature

<b>Technical Proposal ID</b>	<b>4045</b>
<b>Change Date</b>	<b>2018-10-29</b>
<b>Builds on Specification</b>	<b>NVM Express 1.3c, TP 4018a, and TP 4003b</b>

### Technical Proposal Author(s)

<b>Name</b>	<b>Company</b>
Mark Carlson	Toshiba
Judy Brock	Samsung
Tom Friend	SKHynix
Stacey Secatch	Seagate
Ross Stenfort	Facebook

This technical proposal is to develop a method for a host to “tag” the created I/O Queues with an NVMSETID hint to indicate that this is one (of possibly many) targeted at that NVM Set, thus creating an association between a given NVM Set and a given I/O Submission Queue.

### Revision History

<b>Revision Date</b>	<b>Change Description</b>
2018-05-15	Initial version
2018-05-22	Group discussion
2018-06-28	J Brock: <ul style="list-style-type: none"><li>• Rewrote Create I/O Submission Queue Dword 12 text.</li><li>• Added new SQ Associations Config Feature ID section.</li><li>• Added SQ Associations in Controller Attributes field in Identify. Controller and in Controller Initialization flow section.</li><li>• Added section 8 model clause section.</li><li>• Defined error status to return for Invalid NVM Set Identifiers.</li></ul>
2018-07-15	<ul style="list-style-type: none"><li>• Incorporated feedback from 07/10/18 subteam meeting</li></ul>
2018-07-17	<ul style="list-style-type: none"><li>• Minor modifications to section 8.NEW text (modified one line, removed one line)</li></ul>
2018-07-23	<ul style="list-style-type: none"><li>• Incorporated feedback from 07/19/18 Tech WG meeting</li></ul>
2018-07-24	<ul style="list-style-type: none"><li>• Removed SQ Associations Config Feature ID</li><li>• Removed references to enabling this capability before using it</li></ul>
2018-08-15	<ul style="list-style-type: none"><li>• Incorporated Phase 3 feedback from Intel; made host operating rules mandatory instead of being recommendations; mandatory for all SQ to be associated with some NVM Set</li></ul>
2018-08-24	<ul style="list-style-type: none"><li>• Incorporated feedback from 08/23/18 Tech WG meeting</li></ul>
2018-10-11	<ul style="list-style-type: none"><li>• Integration</li></ul>
2018-10-29	<ul style="list-style-type: none"><li>• Ratified</li></ul>

## Description of Specification Changes

*Modify a portion of Figure 211 (Identify – Identify Controller Data Structure) as shown below:*

**Figure 211: Identify – Identify Controller Data Structure**

Bytes	O/M	Description
<b>Controller Capabilities and Features</b>		
01:00	M	<b>PCI Vendor ID (VID):</b> Contains the company vendor identifier that is assigned by the PCI SIG. This is the same value as reported in the ID register in section <b>Error! Reference source not found..</b>
03:02	M	<b>PCI Subsystem Vendor ID (SSVID):</b> Contains the company vendor identifier that is assigned by the PCI SIG for the subsystem. This is the same value as reported in the SS register in section <b>Error! Reference source not found..</b>
...		...
99:96	M	<p><b>Controller Attributes (CTRATT):</b> This field indicates attributes of the controller.</p> <p>Bits 31:9 are reserved.</p> <p><b>Bit 8 (SQ Associations):</b> If set to '1', then the controller supports SQ Associations (refer to section <b>Error! Reference source not found.</b>). If cleared to '0', then the controller does not support SQ Associations.</p> <p><b>Bit 5 (Predictable Latency Mode):</b> If set to '1', then the controller supports Predictable Latency Mode (refer to section <b>Error! Reference source not found.18</b>). If cleared to '0', then the controller does not support Predictable Latency Mode.</p> <p><b>Bit 4 (Endurance Groups):</b> If set to '1', then the controller supports Endurance Groups (refer to section <b>Error! Reference source not found.</b>). If cleared to '0', then the controller does not support Endurance Groups.</p> <p>...</p>
...		...

**Modify a portion of Section 5.4 as shown below:**

#### **5.4 Create I/O Submission Queue command**

The Create I/O Submission Queue command uses the PRP Entry 1, Command Dword 10, ~~and~~ Command Dword 11, ~~and~~ Command Dword 12 fields. All other command specific fields are reserved.

**Figure TBD: Create I/O Submission Queue – Command Dword 12**

Bit	Description
31:16	Reserved
15:00	<p><b>NVM Set Identifier (NVMSETID):</b> This field indicates the identifier of the NVM Set to be associated with this Submission Queue.</p> <p>If this field is cleared to 0h or the SQ Associations capability is not supported (refer to section 8.<b>NEW</b>), then this Submission Queue is not associated with any specific NVM Set.</p> <p>If this field is set to a non-zero value that is not specified in the NVM Set List (refer to figure Fig5_15TBD0) and the SQ Associations feature is supported (refer to section 8.<b>NEW</b>), then the controller shall abort the command with a status code of Invalid Field in Command.</p> <p>The host should not submit commands for namespaces associated with other NVM Sets in this Submission Queue. (refer to section 8.<b>NEW</b>).</p>

**Add new section 8.**NEW**, as shown below:**

#### **8.**NEW** SQ Associations (Optional)**

When Predictable Latency Mode is enabled, all I/O commands for namespaces in a given NVM Set have the same quality of service attributes and shall exhibit predictable latencies as described in section 8.TBD.

The SQ Associations capability provides hints to the controller as to which specific I/O Queues are associated with a given NVM Set. The controller uses this information to further enhance performance when Predictable Latency Mode is enabled.

The SQ Associations capability is an optional capability. Predictable Latency Mode (refer to section 8.TBD) is not dependent on the use of the SQ Associations capability.

If a controller supports SQ Associations, then the controller shall:

- Indicate support for the SQ Associations capability in the Controller Attributes (CTRATT) field in the Identify Controller data structure;
- Indicate support for NVM Sets in the Controller Attributes (CTRATT) field in the Identify Controller data structure; and
- Indicate support for Predictable Latency Mode in the Controller Attributes (CTRATT) field in the Identify Controller data structure.

The host enables the SQ Associations capability by creating an association between an NVM Set and a Submission Queue.

In order for the SQ Associations capability to yield benefits, the host is required to:

- a) create an association between each Submission Queue and some NVM Set; and
- b) only issue I/O commands to Submission Queues that have an association with the NVM Set that contains the namespace associated with the Namespace Identifier specified in that I/O command.

While this capability is enabled, failure to follow the specified operating rules may impact Predictable Latency (refer to section 8.TBD).