



**LEGAL NOTICE:**

© Copyright 2008 to 2022 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 2022 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 Workgroup  
c/o VTM, Inc.  
3855 SW 153<sup>rd</sup> Drive  
Beaverton, OR 97003  
USA  
[info@nvmexpress.org](mailto:info@nvmexpress.org)

## NVM Express™ Technical Proposal (TP)

<b>Technical Proposal ID</b>	<b>4148 Enhanced Namespace Preferred Deallocation Alignment</b>
<b>Change Date</b>	<b>2022-11-30</b>
<b>Builds on Specification(s)</b>	<b>NVM Command Set Specification 1.0b</b>
<b>References</b>	<b>TP4090 Enhanced Deallocation Granularity TP4116 Optimal Read Size and Granularity</b>

### Technical Proposal Author(s)

<b>Name</b>	<b>Company</b>
Judy Brock	Samsung

### Technical Proposal Overview

This proposal will allow a device to advertise a preferred deallocation alignment value larger than is presently possible. Specifically, it will add support for a new 32-bit Namespace Preferred Deallocation Alignment Large (NPDAL) field to address the needs of larger granularities.

### Revision History

<b>Revision Date</b>	<b>Change Description</b>
2022-06-10	Initial version
2022-06-17	Minor editorial changes from WG feedback; accepted all changes and submitted as Phase 2 exit candidate
2022-07-13	30 day review candidate – changed filename to “phase 3”, no known changes to be made in Phase 3. If WG approves start of 30 day review will re-save as “member review”
2022-07-19	Changed one occurrence of “contains” to “is set to”. New 30 day member review candidate (minus file renaming).
2022-11-27	Integrated
2022-11-30	Minor editorial changes from Judy Brock

### Description for Changes Document for NVM Express Command Set 1.0b

#### New Features/Feature Enhancements:

- Added new Namespace Preferred Deallocate Alignment Large (NPDAL) field to the NVM Command Set specific Identify Namespace data structure (CNS 05h, CSI 00h). This is a 32-bit version of the existing NPDA field in the Identify Namespace data structure (CNS 00h).
- Modifies section 5.8.2 to define the usage of the new NPDAL field.

## 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
Purple:	Moved

## Description of Specification Changes for NVM Command Set Specification 1.0b (version dated 2022.04.27)

### 4.1.5.1 NVM Command Set Identify Namespace Data Structure (CNS 00h)

...

#### Figure TBD: Namespace Alignment and Granularity Attributes

<Editor: This Figure is labelled "Figure TBD" because it comes from ratified TP 4116 where it is so-named >

Optimal Write Performance Value	Field Supported					
	NPWG	NPWA	NPDG	NPDGL	NPDA	NOWS
00b	No	No	No	No	No	No
01b	Yes	Yes	Yes	No	Yes	Yes
10b	Yes	Yes	No	Yes	Yes	Yes
11b	Yes	Yes	Yes	Yes	Yes	Yes
The use of these fields by the host for I/O optimization is described in section 5.10.2.						

Field Supported	Optimal Write Performance Value			
	00b	01b	10b	11b
NPWG	No	Yes	Yes	Yes
NPWA	No	Yes	Yes	Yes
NPDG	No	Yes	No	Yes
NPDA	No	Yes	Yes	Yes
NPDGL	No	No	Yes	Yes
NPDAL	No	No	Yes	Yes
NOWS	No	Yes	Yes	Yes
The use of these fields by the host for I/O optimization is described in section 5.8.2.				

### 4.1.5.3 I/O Command Set Specific Identify Namespace Data Structure (CNS 05h)

...

**Figure 100: NVM Command Set I/O Command Set Specific Identify Namespace Data Structure (CSI 00h)**

Bytes	O/M 1	Description	Reported <sup>2</sup>
...			
271:268	O	<p><b>Namespace Preferred Deallocate Granularity Large (NPDGL):</b> This field indicates the recommended granularity in logical blocks for the Dataset Management command with the Attribute – Deallocate bit set to ‘1’ in Command Dword 11. If this field is not supported as defined by the OPTPERF field (refer to Figure 97), then this field is reserved.</p> <p>If this field is cleared to 0h, then this field does not indicate a recommended granularity.</p> <p>The value of this field may change if the namespace is reformatted.</p> <p>If the Namespace Preferred Deallocate Alignment Large (NPDAL) field is cleared to 0h, then the size-value of the NPDGL field should be a multiple of the Namespace Preferred Deallocate Alignment (NPDA) field (refer to Figure 108).</p> <p>If the Namespace Preferred Deallocate Alignment Large (NPDAL) field is supported as defined by the OPTPERF field (refer to Figure 97) and is set to a non-zero value, then the value of the NPDGL field should be a multiple of the NPDAL field.</p> <p>Refer to section 5.8.2 for how this field is utilized to improve performance and endurance.</p>	No
...			
283:280	O	<p><b>Namespace Optimal Read Size (NORS):</b> This field indicates the size in logical blocks for optimal read performance for this namespace. This is a 0's based value. If this field is not supported as indicated by the OPTPERF field, then this field is reserved.</p> <p>The size indicated should be less than or equal to Maximum Data Transfer Size (MDTS) that is specified in units of minimum memory page size. The value of this field may change if the namespace is reformatted. The value of this field should be a multiple of Namespace Preferred Read Granularity (NPRG).</p> <p>Refer to section 5.8.2 for how this field is utilized to improve performance and endurance.</p>	No
287:284	O	<p><b>Namespace Preferred Deallocate Alignment Large (NPDAL):</b> This field indicates the recommended alignment in logical blocks for the Dataset Management command with the Attribute – Deallocate bit set to ‘1’ in Dword 11. If this field is not supported as indicated by the OPTPERF field, then this field is reserved.</p> <p>The value of this field may change if the namespace is reformatted.</p> <p>Refer to section 5.8.2 for how this field is utilized to improve performance and endurance.</p>	No
4095:284288	O	Reserved	
NOTES: 1. O/M definition: O = Optional, M = Mandatory. 2. Identifies fields that report information for the Identify command when querying the capabilities of LBA formats.			

## 5.8.2 Improving Performance through I/O Size and Alignment Adherence

...

The namespace preferred deallocate granularity is a number of logical blocks that is indicated by both the NPDG field (refer to Figure 97) and the NPDGL field. The NPDGL field is able to represent larger values than the NPDG field (refer to Figure 100). Support for these fields is indicated by the OPTPERF field (refer to Figure 97). If the NPDG field and the NPDGL field are both supported and indicate different values of namespace preferred deallocate granularity, then the host should use the value indicated by the NPDGL field.

The namespace preferred deallocate alignment is a number of logical blocks that is indicated by both the NPDA field (refer to Figure 97) and the NPDAL field (refer to Figure 100). The NPDAL field is able to represent larger values than the NPDA field. Support for these fields is indicated by the OPTPERF field (refer to Figure 97). If the NPDA field and the NPDAL field are both supported and indicate different non-zero values of namespace preferred deallocate alignment, then the host should use the value indicated by the NPDAL field. The NPDA field is a 0-based number (i.e. at least one is required to be defined) and the NPDAL field is a 1-based number (i.e., a value of 0h indicates not supported).

...

### 5.8.2.TBD2 Alignment and Read Performance

...

Non-adherence to write-related performance attributes (i.e. NPWG, NPWA, NPDG, NPDGL, NPDA, NPDAL, and NOWS), across all the namespaces in:

- a) the same NVM Set;
- b) the same Endurance Group when NVM Sets are not supported; or
- c) the NVM subsystem when Endurance Groups are not supported,

may affect the level of read optimization achievable through the usage of NORS as described in this section.