



LEGAL NOTICE:

© **Copyright 2007 to 2021 NVM Express™, Inc. ALL RIGHTS RESERVED.**

This NVM Express revision 2.0 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 2.0 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 to 2021 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 153rd Drive
Beaverton, OR 97003
USA
info@nvmexpress.org

NVM Express Technical Proposal for New Feature

Technical Proposal ID	4103 - Controller Optimal Aggregated Queue Depth Reporting
Change Date	2021-06-24
Builds on Specification	NVM Express 2.0
Refers to Ratified Technical Proposals	

Technical Proposal Author(s)

Name	Company
Yoni Shternhell	Western Digital, Inc.

Revision History

Revision Date	Change Description
2021-01-20	Initial version
2021-01-28	Changed in the OAQD field definition
2021-02-15	Clean ups for Phase 3 discussions
2021-02-25	Editorial changes in the OAQD field definition before entering to 30day member review
2021-06-08	Aligned with NVMe 2.0 specification
2021-06-14	Integrated into the NVMe Base Specification, revision 2.0.
2021-06-16	Added Admin and Disc column in the Identify Controller data structure
2021-06-20	Integrated into the NVMe Base Specification, revision 2.0.
2021-06-24	Clean version for integration

Description for NVMe 2.0 Changes Document

This proposal allows an NVMe device to report the optimal aggregated I/O Queue Depth information to enable a host system to dynamically optimize its workload distribution to NVMe devices based on these reported information.

Description of Specification Changes

1. Identify Controller Data Structure extended to include the optimal aggregated I/O Queue Depth information

Markup Conventions:

Black:	Unchanged (however, hot links are removed)
Red Strikethrough:	Deleted
Blue:	New
Highlighted:	TBD values, anchors, and links to be inserted.
Orange Bracketed:	Notes to editor

Modify a portions of Figure 275 (Identify – Identify Controller Data Structure) as shown below:

5.17.2.1 Identify Controller data structure (CNS 01h)

...

Figure 275: Identify – Identify Controller Data Structure

Bytes	I/O ¹	Admin ¹	Disc ¹	Description
...				
NVM Command Set Attributes				
...				
563:560	O	R	R	Maximum I/O Controller Namespace Attachments (MAXCNA): Indicates the maximum number of namespaces that are allowed to be attached to this I/O controller. If this field is cleared to 0h, then no maximum is specified. The value of this field shall be less than or equal to the number of namespaces supported by the NVM Subsystem (refer to the MNAN field).
567:564	O	R	R	Optimal Aggregated Queue Depth (OAQD): Indicates the recommended maximum total number of outstanding I/O commands across all I/O queues on the controller for optimal operation. The host may use this value to limit the number of commands outstanding at one time across all I/O queues on the controller. If this field is cleared to 0h, then the Optimal Aggregated Queue Depth is not reported.
767:568				Reserved

Figure 275: Identify – Identify Controller Data Structure

Bytes	I/O ¹	Admin ¹	Disc ¹	Description
1023:768	M	M	R	<p>NVM Subsystem NVMe Qualified Name (SUBNQN): This field specifies the NVM Subsystem NVMe Qualified Name as a UTF-8 null-terminated string. Refer to section 7.9 for the definition of NVMe Qualified Name.</p> <p>Support for this field is mandatory if the controller supports revision 1.2.1 or later as indicated in the Version register (refer to section 3.1.2).</p>
...				