



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

Technical Proposal ID	TP 4028a
Change Date	11/13/2018
Builds on Specification	NVM Express 1.3c

Technical Proposal Author(s)

Name	Company
Fred Knight	NetApp
David Black	Dell EMC
Sagi Grimberg	LightBits Labs

This technical proposal defines additional common error information related to fabrics and/or pathing related errors.

Several fields were defined for use by TCP that were not needed and never used. This update removes that material and creates a generic transient error code for use by transient transport errors.

Revision History

Revision Date	Change Description
10/18/17	Creation from the ANA (4004) TP
11/20/2017	Add TP # and Authors
12/02/2017	Include just common new error stuff – TP specific stuff left to each specific TP (ANA=TP4004, and TCP=TP4001)
12/13/2017	TP4001 changed its name to TP8000
12/26/2017	Change “Path Related Errors” to “Path Related Status”
12/29/2017	Resolve comments from Toshiba, Add PCI transport type questions (to be resolved at full group meeting along with RDMA question); make transport type specific field more generic for possible use by other future transports.
1/4/2018	Updates from the group call; update definition of 3h status code.
5/23/2018	Move transport specific coded values to the Transport Specs. Therefore, also removed error code values that had been reserved for TCP and replace them with generic errors. This improves layering separation between the NVMe layer and the transport specific driver layers. Clarify that Pathing errors are for ANA.
9/10/2018	Define the transient transport error code.
10/30/2018	Integration
11/13/2018	Ratified

Description of Specification Changes

Make changes to section 4.6.1.1 (Status Code Type) as shown below:

4.6.1.1 Status Code Type (SCT)

Completion queue entries indicate a status code type for the type of completion being reported. Figure 30 specifies the status code type values and descriptions.

Figure 1: Status Code – Status Code Type Values

Value	Description	Reference
0h	Generic Command Status: Indicates that the command specified by the Command and Submission Queue identifiers in the completion queue entry has completed. These status values are generic across all command types, and include such conditions as success, opcode not supported, and invalid field.	4.6.1.2.1
1h	Command Specific Status: Indicates a status value that is specific to a particular command opcode. These values may indicate additional processing is required. Status values such as invalid firmware image or exceeded maximum number of queues is reported with this type.	4.6.1.2.2
2h	Media and Data Integrity Errors: Any media specific errors that occur in the NVM or data integrity type errors shall be of this type.	4.6.1.2.3
3h	Path Related Status: Indicates that the command specified by the Command and Submission Queue identifier in the completion queue entry has completed. These status values are generic across all command types. These values may indicate that additional process is required and indicate a status value that is specific to: <ul style="list-style-type: none"> a) the connection between the host and the controller processing the command; or b) for controllers that support Asymmetric Namespace Access Reporting (refer to section 8.20), the characteristics of the relationship between the controller processing the command and the specified namespace. 	4.6.1.2.4
43 h – 6h	Reserved	
7h	Vendor Specific	

Make changes to section 4.6.1.2.1 (Generic Command Status Definition) as shown below:

4.6.1.2.1 Generic Command Status Definition

Completion queue entries with a Status Code type of Generic Command Status indicate a status value associated with the command that is generic across many different types of commands.

Figure 31: Status Code – Generic Command Status Values

Value	Description
00h	Successful Completion: The command completed without error.
...	...
21h	...
22h	Transient Transport Error: A transient transport error was detected. If the command is retried on the same controller, the command is likely to succeed. A command that fails with a transient transport error four or more times should be treated as a persistent transport error that will not succeed if retried on the same controller.
2023 h – 7Fh	Reserved
80h – BFh	I/O Command Set Specific
C0h – FFh	Vendor Specific

Add new section 4.6.1.2.4 (Path Related Status Definition) as shown below:

4.6.1.2.4 Path Related Status Definition

Completion queue entries with a Status Code type of Path Related Status (refer to Figure 36a) indicate a status value associated with the command that is generic across many different types of commands and applies to a specific connection between the host and controller processing the command or between the controller and the namespace. The command for which this status is returned may be retried on a different controller in the same NVM subsystem if more than one controller is available to the host.

In a multipath environment, unless otherwise specified, errors of this type should be retried using a different path, if one is available.

Figure 36a: Status Code – Path Related Status Values

Value	Description
00h	Internal Path Error: The command was not completed as the result of a controller internal error that is specific to the controller processing the command. Retries for the requested function should be based on the setting of the DNR bit (refer to Figure 29).
01h-03h	Reserved for TP4004
04h-5Fh	Reserved
Controller detected Pathing errors	
60h	Controller Pathing Error: A pathing error was detected by the controller.
61h-6Fh	Reserved
Host detected Pathing errors	
70h	Host Pathing Error: A pathing error was detected by the host.
71h	Command Aborted By host: The command was aborted as a result of host action (e.g., the host disconnected the Fabric connection).
72h-7Fh	Reserved
80h – BFh	I/O Command Set Specific
C0h – FFh	Vendor Specific

Make changes to section Figure 92 (Get Log Page – Error Information Log Entry) as shown below:

5.14.1.1 Error Information (Log Identifier 01h)

...

Figure 32: Get Log Page – Error Information Log Entry (Log Identifier 01h)

Bytes	Description
...	...
27:24	Namespace: This field indicates the NSID of the namespace that the error is associated with, if applicable.
...	...
63:40	Reserved