



**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 153<sup>rd</sup> Drive  
Beaverton, OR 97003 USA  
info@nvmexpress.org

## NVM Express Technical Proposal for New Feature

Technical Proposal ID	TP 4028
Change Date	01/09/2018
Builds on Specification	NVM Express 1.3

### 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.

### 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.
1/18/2018	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	<b>Generic Command Status:</b> 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	<b>Command Specific Status:</b> 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	<b>Media and Data Integrity Errors:</b> 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	<b>Path Related Status:</b> Indicates a status value that is specific to: a) the connection between the host and the controller processing the command; or b) the characteristics of the relationship between the controller processing the command and the specified namespace.	4.6.1.2.4
4h – 6h	Reserved	
7h	<b>Vendor Specific</b>	

**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 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	<b>Internal Path Error:</b> 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 request function should be based on the setting of the DNR bit (refer to Figure 29).
01h-03h	<b>Reserved for TP4004</b>
04h-5Fh	<b>Reserved</b>
Controller detected Pathing errors	
60h	<b>Controller Pathing Error:</b> A pathing error was detected by the controller.
61h-63h	<b>Reserved for TP8000</b>
64h-6Fh	<b>Reserved</b>
Host detected Pathing errors	
70h	<b>Host Pathing Error:</b> A pathing error was detected by the host.

Value	Description
71h	<b>Command Aborted By host:</b> The command was aborted as a result of host action (e.g., the host disconnected the Fabric connection).
72h-74h	<b>Reserved for TP8000</b>
75h-7Fh	<b>Reserved</b>
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 2: Get Log Page – Error Information Log Entry (Log Identifier 01h)**

Bytes	Description								
07:00	<b>Error Count:</b> This is a 64-bit incrementing error count, indicating a unique identifier for this error. The error count starts at 1h, is incremented for each unique error log entry, and is retained across power off conditions. A value of 0h indicates an invalid entry; this value is used when there are lost entries or when there are fewer errors than the maximum number of entries the controller supports.								
09:08	<b>Submission Queue ID:</b> This field indicates the Submission Queue Identifier of the command that the error information is associated with. If the error is not specific to a particular command then this field shall be set to FFFFh.								
11:10	<b>Command ID:</b> This field indicates the Command Identifier of the command that the error is associated with. If the error is not specific to a particular command then this field shall be set to FFFFh.								
13:12	<b>Status Field:</b> This field indicates the Status Field for the command that completed. The Status Field is located in bits 15:01, bit 00 corresponds to the Phase Tag posted for the command. If the error is not specific to a particular command then this field reports the most applicable status value.								
15:14	<p><b>Parameter Error Location:</b> This field indicates the byte and bit of the command parameter that the error is associated with, if applicable. If the parameter spans multiple bytes or bits, then the location indicates the first byte and bit of the parameter.</p> <table border="1"> <thead> <tr> <th>Bits</th><th>Description</th></tr> </thead> <tbody> <tr> <td>15:11</td><td>Reserved</td></tr> <tr> <td>10:8</td><td>Bit in command that contained the error. Valid values are 0 to 7.</td></tr> <tr> <td>7:0</td><td>Byte in command that contained the error. Valid values are 0 to 63.</td></tr> </tbody> </table> <p>If the error is not specific to a particular command then this field shall be set to FFFFh.</p>	Bits	Description	15:11	Reserved	10:8	Bit in command that contained the error. Valid values are 0 to 7.	7:0	Byte in command that contained the error. Valid values are 0 to 63.
Bits	Description								
15:11	Reserved								
10:8	Bit in command that contained the error. Valid values are 0 to 7.								
7:0	Byte in command that contained the error. Valid values are 0 to 63.								
23:16	<b>LBA:</b> This field indicates the first LBA that experienced the error condition, if applicable.								
27:24	<b>Namespace:</b> This field indicates the <b>NSID of the</b> namespace that the error is associated with, if applicable.								
28	<b>Vendor Specific Information Available:</b> If there is additional vendor specific error information available, this field provides the log page identifier associated with that page. A value of 00h indicates that no additional information is available. Valid values are in the range of 80h to FFh.								

29	<p><b>Transport Type (TRTYPE):</b> This field indicates the Transport Type of the transport associated with the error. The values in this field are the same as the TRTYPE values in the Discovery Log Page Entry (refer to the NVMe over Fabrics specification). If the error is not transport related, this field shall be cleared to '0'. If the error is transport related, this field shall be set to the type of the transport as follows:</p> <table> <tr> <th>Value</th><th>Description</th></tr> <tr> <td>0h</td><td>The transport type is not indicated or the error is not transport related.</td></tr> <tr> <td>1h</td><td>RDMA Transport (refer to the NVMe over Fabrics specification)</td></tr> <tr> <td>2h</td><td>Fibre Channel Transport (refer to INCITS 540)</td></tr> <tr> <td>3h</td><td>TCP Transport (refer to TBD in TP8000)</td></tr> <tr> <td>FEh</td><td>Intra-host Transport (i.e., loopback) (Note: This is a reserved value for use by host software.)</td></tr> <tr> <td>All others</td><td>Reserved</td></tr> </table>	Value	Description	0h	The transport type is not indicated or the error is not transport related.	1h	RDMA Transport (refer to the NVMe over Fabrics specification)	2h	Fibre Channel Transport (refer to INCITS 540)	3h	TCP Transport (refer to TBD in TP8000)	FEh	Intra-host Transport (i.e., loopback) (Note: This is a reserved value for use by host software.)	All others	Reserved
Value	Description														
0h	The transport type is not indicated or the error is not transport related.														
1h	RDMA Transport (refer to the NVMe over Fabrics specification)														
2h	Fibre Channel Transport (refer to INCITS 540)														
3h	TCP Transport (refer to TBD in TP8000)														
FEh	Intra-host Transport (i.e., loopback) (Note: This is a reserved value for use by host software.)														
All others	Reserved														
31: <del>29</del> 30	Reserved														
39:32	<b>Command Specific Information:</b> This field contains command specific information. If used, the command definition specifies the information returned.														
41:40	<p><b>Transport Type Specific Information:</b> This field indicates additional transport type specific error information. If multiple errors exist, then this field indicates additional information about the first error. This field is transport type dependant (refer to TRTYPE) as follows:</p> <table> <tr> <th>Transport Type</th><th>Description</th></tr> <tr> <td>All others</td><td>Reserved</td></tr> <tr> <td>3h</td><td>This field indicates, the offset, in bytes, from the start of the Transport Header to the start of the field that is in error. If multiple errors exist, then this field indicates the lowest offset that is in error.</td></tr> </table>	Transport Type	Description	All others	Reserved	3h	This field indicates, the offset, in bytes, from the start of the Transport Header to the start of the field that is in error. If multiple errors exist, then this field indicates the lowest offset that is in error.								
Transport Type	Description														
All others	Reserved														
3h	This field indicates, the offset, in bytes, from the start of the Transport Header to the start of the field that is in error. If multiple errors exist, then this field indicates the lowest offset that is in error.														
63: <del>40</del> 42	Reserved														