



#### **LEGAL NOTICE:**

© **Copyright 2007 - 2018 NVM Express, Inc. ALL RIGHTS RESERVED.**

This NVM Express over Fabrics revision 1.0 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 over Fabrics revision 1.0 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 - 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 service marks may be claimed as the property of their respective owners.

NVM Express Workgroup  
c/o VTM, Inc.  
3855 SW 153rd Drive  
Beaverton, OR 97003  
info@nvmexpress.org

## NVM Express Technical Proposal for New Feature

Technical Proposal ID	NVMe over Fabrics 1.0 TP 8008
Change Date	11/16/2016
Builds on Specification	NVM Fabrics 1.0

### Technical Proposal Author(s)

Name	Company
Idan Burstein	Mellanox
Dave Minturn	Intel
Parav Pandit	HPE
Christoph Hellwig	WD

The proposal defines the Controller ID field to be passed within the RDMA Queue Pair creation for those associated with I/O Queues to enable optimizations such as allocating the number of entries in the queue appropriately.

There are transport specific status codes that need to be reported in a standard fashion. A range of status codes has been allocated based on a request from T11 for Fibre Channel.

### Revision History

Revision Date	Change Description
08/25/2016	Removed the CPTH and TRBA, added the Controller ID to the RDMA_CM
8/31/2016	Editorial updates. Added status code reservation based on Fibre Channel needs. Added UUID for software based endpoints as an alternative to NGUID.
9/15/2016	Updates for UUID identifier.
9/25/2016	Final updates from Technical team. Pushed namespace UUID identifier to a base specification Technical Proposal.
11/16/2016	Ratified.

## Description of Specification Changes

**Modify NVMe over Fabrics 1.0 Figure 44 as shown:**

**Figure 44: RDMA\_CM\_REQUEST Private Data Format**

Byte	Description
01:00	<b>Record Format (RECFMT):</b> Specifies the format of the RDMA Private Data. If a new format is defined, this value is incremented by one. The format of the record specified in this definition shall be 0h.
03:02	<b>Queue ID (QID):</b> Refer to the Queue ID definition in the Connect command in Figure 19.
05:04	<b>RDMA QP Host Receive Queue Size (HRQSIZE):</b> This field indicates the number of RDMA QP receive queue entries allocated by the host's RDMA Transport for capsule reception.
07:06	<b>RDMA QP Host Send Queue Size (HSQSIZE):</b> This field indicates the number of RDMA QP send queue entries allocated by the host's RDMA transport for capsule transmission. The value shall be set to the Submission Queue Size (SQSIZE). Refer to the SQSIZE definition in the Connect command in Figure 19.
09:08	<b>Controller ID (CNTLID):</b> This field is used to indicate the Controller ID if the RDMA QP is for an I/O Queue. If the QID field is cleared to 0h (i.e., the RDMA QP is for an Admin Queue), this field shall be cleared to 0h. If the QID field is non-zero, this field should be set to the value of the Controller ID associated with the creation of this queue.
31:10	Reserved

**Figure 45: RDMA\_CM\_ACCEPT Private Data Format**

Byte	Description
01:00	<b>Record Format (RECFMT):</b> Specifies the format of the RDMA Private Data. If a new format is defined, this value is incremented by one. The format of the record specified in this definition shall be 0h.
03:02	<b>RDMA QP Controller Receive Queue Size (CRQSIZE):</b> This field indicates the number of RDMA QP receive queue entries allocated by the controller's RDMA Transport for capsule reception. RDMA Transports that use RNR_RETRY flow control may set this entry to be less than or equal to the value of HSQSIZE specified in Figure 43. RDMA Transports that do not use RNR_RETRY shall set this value to be equal to the value of HSQSIZE specified in Figure 43.
31:04	Reserved

**Figure 46: RDMA\_CM\_REJECT Private Data Format**

Byte	Description
01:00	<b>Record Format (RECFMT):</b> Specifies the format of the RDMA Private Data. If a new format is defined, this value is incremented by one. The format of the record specified in this definition shall be 0h.
03:02	<b>Status (STS):</b> Specifies status for the associated RDMA_CM_REQUEST that is paired with this reject response. The valid status values are specified in Figure 46.

**Modify NVMe over Fabrics 1.0 Figure 46 as shown below:**

**Figure 46: RDMA Transport Errors**

Value	Description
1h	<b>RDMA Invalid Private Data Length:</b> The host sent an incorrect private_data size.
2h	<b>RDMA Invalid RECFMT:</b> The host sent an invalid RECFMT.
3h	<b>RDMA Invalid QID:</b> The host sent an invalid QID.
4h	<b>RDMA Invalid HSQSIZE:</b> The host sent an invalid HSQSIZE.
5h	<b>RDMA Invalid HRQSIZE:</b> The host sent an invalid HRQSIZE.
6h	<b>RDMA No Resources:</b> The controller-side RDMA transport is unable to create the RDMA QP due to lack of resources.
7h	<b>RDMA Invalid IRD:</b> The host sent an invalid IRD value.
8h	<b>RDMA Invalid ORD:</b> The host sent an invalid ORD value.
9h	<b>RDMA Invalid CNTLID:</b> The host sent an invalid CNTLID value.
Ah – FFh	Reserved

**Modify NVMe over Fabrics Figure 10 as shown below:**

**Figure 10: Fabrics Command Specific Status Values**

Value	Description	Commands Affected
80h	<b>Connect Incompatible Format:</b> The NVM subsystem does not support the Connect Record Format specified by the host.	Connect
81h	<b>Connect Controller Busy:</b> The controller is already associated with a host. This value is also returned if there is no available controller.	Connect
82h	<b>Connect Invalid Parameters:</b> One or more of the parameters (Host NQN, Subsystem NQN, Host Identifier, Controller ID, Queue ID) specified are not valid.	Connect
83h	<b>Connect Restart Discovery:</b> The NVM subsystem requested is not available. The host should restart the discovery process.	Connect
84h	<b>Connect Invalid Host:</b> The host is not allowed to establish an association to any controller in the NVM subsystem or the host is not allowed to establish an association to the specified controller.	Connect
85h – 8Fh	Reserved	
90h	<b>Discover Restart:</b> The snapshot of the records is now invalid or out of date. The host should re-read the Discovery Log Page.	Get Log Page
91h	<b>Authentication Required:</b> NVMe in-band authentication is required and the queue has not yet been authenticated.	NOTE 1
92h – <del>BFh</del> AFh	Reserved	
B0h – BFh	<b>Transport Specific:</b> The status values in this range are NVMe Transport specific. Refer to the appropriate NVMe Transport binding specification for the definition of these status values.	
NOTES: 1. All commands other than Connect, Authenticate Send, and Authenticate Receive.		