



LEGAL NOTICE:

© **Copyright 2008 to 2023 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 2023 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 (TP)

Technical Proposal ID	TP 4175 Copy and Reachability
Revision Date	2023-10-31
Builds on Specifications	NVM Express NVM Command Set Specification 1.0c
References	TP 4130 Cross-Namespace Copy TP 4156 Reachability Architecture

Technical Proposal Author(s)

Name	Company
Fred Knight	NetApp
John Meneghini	RedHat
Murali Rajagopal	VMware

Technical Proposal Overview

Associate the error from the Reachability Architecture (TP4156) in the context of the Copy command (TP4131).

Revision History

Revision Date	Change Description
2023-06-01	Initial rough version
2023-06-06	Incorporate changes from FMDS meeting, remove Base spec content (as it just repeated what was already in TP4156); reword non-reachability error paragraph; relocate text further down in Copy Command section (after SREs are discussed).
2023-06-26	Move duplicate command specific error text into a common paragraph.
2023-07-13	Update description of the change, update the description of the meaning of the error to sort out the difference between NSID and namespace, and add a reference.
2023-08-29	Update for member review
2023-10-24	Incorporate member review comments
2023-10-31	Integrated

Description for Changes Document for NVM Command Set Specification 1.0c

New Features/Feature Enhancements/Required Changes:

- Define the relationship between Reachability and the Copy command.

References:

- NVM Command Set Specification 1.0c
- Technical Proposal 4130: Cross-Namespace Copy
- Technical Proposal 4156: Reachability Architecture

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
Purple	Text moved without change.
Orange	Text copied from a referenced TP.
<Green Bracketed>:	Notes to editor

Description of Specification Changes for NVM Command Set Specification 1.0c:

Modify section 3.2.2 as shown below:

3.2.2 Copy command

...

The number of logical blocks written by the Copy command is the sum of all Number of Logical Blocks fields in all Source Range entries specified in the list of Source Range entries.

The data bytes in the LBAs specified by each Source Range Entry shall be copied to the destination LBA range in the same order those LBAs are listed in the Source Range entries (e.g., the LBAs specified by Source Range entry 0 are copied to the lowest numbered LBAs specified by the SDLBA field, the LBAs specified by Source Range entry 1 are copied to the next consecutively numbered LBAs specified by the SDLBA field). The read operations and write operations used to perform the copy may operate sequentially or in parallel.

Two LBA ranges overlap if they specify LBAs in the same namespace and there is at least one LBA that is part of both LBA ranges.

For Source Range Entries Copy Descriptor Formats 0h and 1h, The host should not specify a destination LBA range that overlaps the LBA in any of the Source Range entries. If the host specifies a destination LBA range that overlaps with any LBAs specified in one or more of the Source Range entries, then upon completion of the Copy command, the data stored in each logical block in that overlapping destination LBA range may, within the constraints of the atomicity rules described in section 2.1.4, be from any of the one or more Source Range entries in which that LBA is contained. This is a result of the possibility that overlapping Source Range entries may be processed in any order.

For Source Range Entries Copy Descriptor Formats 2h and 3h, overlap of any source LBA range that is located in the destination namespace with the destination LBA range is prohibited. If a Copy command uses either Source Range Entries Copy Descriptor Format 2h or 3h and any specified source LBA range that is located in the destination namespace has any LBAs in common with the specified destination LBA range, then the controller shall abort the command with a status code of Overlapping I/O Range.

If the controller:

- supports reachability reporting (refer to the Reachability Reporting architecture section in the NVM Express Base Specification);
- processes a Copy command that requests a copy between different namespaces (i.e., the NSID in a Source Range is different than the NSID of the destination namespace); and
- does not report a Reachability Association (refer to the Reachability Reporting architecture section in the NVMe Base Specification) between those namespaces,

then the controller shall abort the command with the status code set to Namespace Not Reachable.

If the read portion of a copy operation attempts to access a deallocated or unwritten logical block, the controller shall operate as described in section 3.2.3.2.1.

Figure 36 shows an example of the relationship between the source LBAs and the destination LBAs.

...

Modify section 3.2.2.TBD5 <from TP4130> as shown below:

3.2.2.TBD5 Command Completion

...

Copy command specific errors are defined in Figure 37.

Figure 37: Copy – Command Specific Status Values

Value	Description
...	
8Ah	Namespace Not Reachable: One or more of the Source Range Entries specifies the NSID of a namespace that is not contained in any Reachability Group that is in a Reachability Association with the Reachability Group that contains the destination namespace. Refer to the Reachability Reporting architecture section in the NVM Express Base Specification.