



LEGAL NOTICE:

© **Copyright 2008 to 2022 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 2022 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


Technical Proposal ID	4093a – Zone Relative Data Lifetime Hint
Change Date	2022-08-14
Builds on Specification	NVMe - Zoned Namespace Command Set Specification 1.1
Refers to Ratified Technical Proposals	TP4053a

Technical Proposal Author(s)

Name	Company
Yoni Shternhell, Matias Bjørling	Western Digital, Inc.

This proposal will add a method for a host to specify to the device a relative lifetime hint upon opening a zone. The hint communicates the relative lifetime of the data that will be subsequently written to that zone.

The relative lifetime hint is defined as a range (e.g., 1-63), that allows the host to bucketize its data according to the relative lifetime of the data. The range itself is monotonically increasing (greater than or equal to). For example, a relative lifetime value of 1 advertises the shortest lifetime of the data to be written, and a value of 63 specifies the longest lifetime of the data to be written. The device can use the specified value and the monotonicity property to optimize its data placement.



As an example, a host splits the lifetime of the data it writes to zones into eight buckets. The relative lifetime is monotonically increasing from bucket 1 to bucket 8 (i.e., the first bucket is data that has the shortest relative lifetime, and the last bucket is data with the longest relative lifetime). When the host opens a zone, it specifies the relative lifetime of the data that will be subsequently written to the zone.

Another example is a host that splits the lifetime of the data into eight buckets, but initially spreads them across the 63 values. Another host may only split its data into two buckets (e.g., 1 and 63), or simply not specify any value. Ultimately, the device internally maps as appropriate by means that are not part of this proposal.

Revision History

Revision Date	Change Description
2020-10-06	Initial version
2020-11-13	Change the Data Lifetime field to be MSB
2020-12-08	Editorial changes to the Data Lifetime field
2021-01-12	Incorporated comments from Phase 2 exit vote
2021-01-26	Changes in the Data Lifetime field <ul style="list-style-type: none">- Make it clear that the scope is namespace- More editorial changes
2021-02-01	Added some text that the value can only be set once per open
2021-02-02	Change of the bit ordering of the Data Lifetime field and move text from the field's definition to the Open Zone section
2021-02-09	Editorial text changes for the Data Lifetime field
2021-02-15	Cleaning for Phase 3 exit ballot
2021-02-16	<ul style="list-style-type: none">• In the description of the Data Lifetime field, three occurrences of "indicates" were changed to "specifies"• We modified "A value of 1h indicates the shortest estimated period before the zone is reset" to "A value of 1h specifies the shortest Data Lifetime." Similar change is in the 63h value text.• Editorial changed 4.3.1.3
2021-02-17	<ul style="list-style-type: none">• Modified text based on members comments:<ul style="list-style-type: none">○ Fixed navigation pan○ Editorial changes○ Remove the Read-Only state from the list in 4.3.1.3
2021-02-23	<ul style="list-style-type: none">• Removing addressed comments and more editorial cleanup
2021-02-25	<ul style="list-style-type: none">• Editorial cleanups before entering to 30day member review
2021-03-22	<ul style="list-style-type: none">• Members review comments<ul style="list-style-type: none">○ Changed the value of the longest Data Lifetime to 3Fh (from 63h)
2021-06-08	<ul style="list-style-type: none">• Aligned to NVMe 2.0 specification
2021-06-20	<ul style="list-style-type: none">• Integrated into the NVMe Zoned Namespace Command Set Specification, revision 1.1.
2021-06-24	<ul style="list-style-type: none">• Clean version for integration
2022-03-30	<ul style="list-style-type: none">• Replacing 'active state' to ZSEO:Explicitly Opened state or the ZSC:Closed state• Updating Figure 30 with Set Zone Descriptor Extension Zone Send Action
2022-04-27	<ul style="list-style-type: none">• Accepted all changes• Clean version ready for 30-day members review
2022-08-10	<ul style="list-style-type: none">• Integrated
2022-08-14	<ul style="list-style-type: none">• Fixed the heading for section 3.4.3.

Description for NVMe 2.0 Changes Document

Description of Zoned Namespace Command Set Specification Changes

TP 4053a adds base support for Zoned Namespaces. The proposal will add a method for a host to specify to the device a relative lifetime hint upon opening a zone. The hint communicates the relative lifetime of the data that will be subsequently written to that zone.

Description of Specification Changes

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

3.4.3 Zone Management Send command

Figure 39: Zone Management Send – Command Dword 13

Bits	Description				
31:0924	Reserved				
23:16	Zone Management (ZM): This field indicates attributes for the Zone Send Action being issued.				
	Zone Send Action	Description			
		Open Zone	Bits	Attribute	Definition
			07:06	Reserved	
			05:00	Data Lifetime	The value in this field specifies an estimate of the relative Data Lifetime of this zone as compared to the relative Data Lifetime of other zones in this namespace. Refer to section 3.4.3.1.3.
					A value of 0h specifies that no estimate is provided.
A value of 1h specifies the shortest Data Lifetime.					
Set Zone Descriptor Extension			A value of 3Fh specifies the longest Data Lifetime.		
All other values	Reserved				
15:09	Reserved				
08	Select All: If this bit is set to '1', then the SLBA field shall be ignored. If this bit is cleared to '0', then the SLBA field specifies the lowest logical block of the zone. Refer to section 4.3.1 for specific behavior for each Zone Send Action.				

3.4.3.1.3 Open Zone:

...

If there are insufficient available Active Resources or insufficient available Open Resources, then the command shall be aborted as described in section 2.1.1.4 and no zone state transition shall occur.

The Data Lifetime of a zone is relative to the Data Lifetimes of other zones in the same namespace.

The Data Lifetime of a zone is not defined until a Zone Management Send command is processed that transitions that zone from the ZSE:Empty state to the ZSEO:Explicitly Opened state or the ZSC:Closed state.

The Data Lifetime of a zone (refer to Figure 39) represents the expected time from a Zone Management Send command transitioning that zone from the ZSE:Empty state to the ZSEO:Explicitly Opened state or the ZSC:Closed state (refer to section 2.1.1.4) until that zone transitions to:

- a) the ZSE:Empty state; or
- b) the ZSO:Offline state.