



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

Technical input submitted to NVM Express® is subject to the terms of the NVM Express® Participant's agreement. Copyright © 2008 to 2022 NVM Express, Inc.

NVM Express® Technical Proposal (TP)

Technical Proposal ID	4115 - ZNS Namespace Management Enhancements
Revision Date	2022.03.15
Builds on Specification(s)	NVM Express® Zoned Namespace Command Set Specification 1.1a NVM Express® NVM Command Set Specification 1.0a
References	TP 4076 Zoned Random Write Area ECN106

Technical Proposal Author(s)

Name	Company
WDC	Yoni Shternhell, Matias Bjorling

Technical Proposal Overview

This proposal improves the management of zoned namespaces by allowing the host to specify how many zone resources (e.g., active resources, open resources, ZRWA resources) are allocated to a specific namespace when that namespace is created. The intent with this TP is to only define the ability to statically allocated resources to a namespace. Resources are not moveable between namespaces unless namespace are fully deleted and recreated.

Revision History

Revision Date	Change Description
2021.08.13	Initial draft
2021.08.25	Incorporated changes from ZNS TG meeting
2021.09.08	Incorporated changes from ZNS TG meeting
2021.09.15	Modifications to the Namespace Management command section
2021.09.16	<ul style="list-style-type: none"> Updated the MAR, MOR and NUMZRWA fields with the FFh and Feh cases Moved the Total and Unallocated resources from Controller to a Domain scope
2021.10.06	<ul style="list-style-type: none"> Major rewrite to the RAR, ROR and RNUMZRWA fields Change the reference in the Host Software Specified Fields to be more general (and not the to the ZNS CS)
2021.10.13	<ul style="list-style-type: none"> Accepting most of the changes provided in the Oct 6 ZNS TG meeting
2021.10.19	<ul style="list-style-type: none"> Update the “Description of changes” section Moved text from section 4.1.6 into the RAR, ROR and RNUMZRWA fields definition Moved completion criteria text to a new command completion section in the NVMe ZNS CS
2021.10.27	<ul style="list-style-type: none"> Accepted all changes
2021.11.03	<ul style="list-style-type: none"> Phase 3 changes We changed the text in the Total and Unallocated resources fields definition: <ul style="list-style-type: none"> Replaced the word ‘sum’ with ‘total Removed the text “allowed to be allocated to namespaces”
2021.11.17	<ul style="list-style-type: none"> Accepted all editorial changes
2021.12.01	<ul style="list-style-type: none"> Member review comment resolution
2022.01.21	<ul style="list-style-type: none"> Aligned with ECN106 text.
2022.03.06	<ul style="list-style-type: none"> Integrated
2022.03.09	<ul style="list-style-type: none"> Resolve editorial integration comments
2022.03.10	<ul style="list-style-type: none"> Change “this field is ignored” wording to “field should be ignored by the host”
2022.03.14	<ul style="list-style-type: none"> Corrected text color and font issues
2022.03.15	<ul style="list-style-type: none"> Fixed the name of the company in the footer. Updated section 4.1.6 to have just one list. Corrected the table sub headings in Figure TBDX.

Description for Changes Document for NVM Express® NVM Command Set Specification 1.0a

New Features/Feature Enhancements/Required Changes:

- Namespace Management – Host Software Specified Fields
 - Description of change.
 - Reserved space for 3 new fields: Requested Active Resources (RAR), Requested Open Resources (ROR), and Requested Number of ZRWA Resources Management (RNUMZRWA)
 - References
 - Technical Proposal 4115

Description for Changes Document for NVM Express® Zoned Namespace Command Set Specification 1.1a

New Features/Feature Enhancements/Required Changes:

- I/O Command Set specific Identify Controller data structure
 - Description of change.
 - 3 new resources fields for total values: Total Active Resources, Total Open Resources and Total ZRWA Resources
 - 3 new resources fields for Unallocated values: Unallocated Active Resources, Unallocated Open Resources and Unallocated ZRWA Resources
 - New Identify field (Namespace Resource Management) in a new attribute (Zoned Controller Attributes)
 - New data structure in the Namespace Management command for Host Software Specified Fields
 - References
 - Technical Proposal 4115

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.
Orange:	Text from ECN106.
<Green Bracketed>:	Notes to editor

Description of Specification Changes for NVM Express® NVM Command Set Specification 1.0a

4 Admin Command Set

...

4.1.6 Namespace Management command

...

Figure 105: Namespace Management – Host Software Specified Fields

Bytes	Description	Host Specified
Fields that are a subset of the Identify Namespace data structure (refer to Figure 97)		
07:00	Namespace Size (NSZE)	Yes
15:08	Namespace Capacity (NCAP)	Yes
25:16	Reserved	
26	Formatted LBA Size (FLBAS)	Yes
28:27	Reserved	
29	End-to-end Data Protection Type Settings (DPS)	Yes
30	Namespace Multi-path I/O and Namespace Sharing Capabilities (NMIC)	Yes
91:31	Reserved	
95:92	ANA Group Identifier (ANAGRPID) ¹	Yes
99:96	Reserved	
101:100	NVM Set Identifier (NVMSETID) ¹	Yes
103:102	Endurance Group Identifier (ENDGID)	Yes
383:104	Reserved	
Fields that are not a subset of the Identify Namespace data structure.		
391:384	Logical Block Storage Tag Mask (LBSTM)	Yes
498:392	Reserved	
511:499	Reserved for I/O Command Sets that extend this specification. Refer to the applicable I/O Command Set specification (e.g., Zoned Namespace Command Set Specification).	
Notes:		
1. A value of 0h specifies that the controller determines the value to use (refer to section 8.12). If the associated feature is not supported, then this field is ignored by the controller.		

Description of Specification Changes for Zoned Namespace Command Set Specification 1.1a

4. Admin Commands for the Zoned Namespace Command Set

4.1 Admin Command behavior for the Zoned Namespace Command Set

4.1.5 Identify Command

4.1.5.2 I/O Command Set Specific Identify Controller Data Structure (CNS 06h, CSI 02h)

Figure 50 defines the I/O Command Set Specific Identify Controller data structure for the Zoned Namespace Command Set.

Figure 50: I/O Command Set Specific Identify Controller Data Structure for the Zoned Namespace Command Set

Bytes	O/M ¹	Description								
00	O ²	<p>Zone Append Size Limit (ZASL): If the Zone Append command is supported then:</p> <ul style="list-style-type: none">a) a non-zero value in this field indicates the maximum data transfer size for the Zone Append command (refer to section 3.4.1); andb) a value of 0h in this field indicates that the maximum data transfer size for the Zone Append command is indicated by the Maximum Data Transfer Size (MDTS) field (refer to the NVMe Base Specification). <p>The value is in units of the minimum memory page size (CAP.MPSMIN) and is reported as a power of two (2ⁿ). This field includes the length of metadata if metadata is interleaved with the stored logical block data.</p> <p>The value of this field shall be less than or equal to the Maximum Data Transfer Size (MDTS).</p>								
03:01		Reserved								
07:04	O ³	<table><tr><td colspan="2">Zoned Controller Attributes (ZCTRATT): This field indicates attributes of the controller that support this command set.</td></tr><tr><th>Bits</th><th>Description</th></tr><tr><td>31:01</td><td>Reserved</td></tr><tr><td>00</td><td>Zoned Namespace Resource Management (ZNRM): If set to '1', then the controller supports allocation of resources when creating a zoned namespace (refer to section 4.1.6). If cleared to '0', then the controller does not support allocation of resources when creating a zoned namespace.</td></tr></table>	Zoned Controller Attributes (ZCTRATT): This field indicates attributes of the controller that support this command set.		Bits	Description	31:01	Reserved	00	Zoned Namespace Resource Management (ZNRM): If set to '1', then the controller supports allocation of resources when creating a zoned namespace (refer to section 4.1.6). If cleared to '0', then the controller does not support allocation of resources when creating a zoned namespace.
Zoned Controller Attributes (ZCTRATT): This field indicates attributes of the controller that support this command set.										
Bits	Description									
31:01	Reserved									
00	Zoned Namespace Resource Management (ZNRM): If set to '1', then the controller supports allocation of resources when creating a zoned namespace (refer to section 4.1.6). If cleared to '0', then the controller does not support allocation of resources when creating a zoned namespace.									
11:08	O ³	Total Active Resources (TAR): This field indicates the total number of active resources (i.e., allocated and unallocated active resources). If the ZNRM bit is cleared to '0', then this field should be ignored by the host.								
15:12	O ³	Unallocated Active Resources (UAR): This field indicates the number of unallocated active resources. If the ZNRM bit is cleared to '0', then this field should be ignored by the host.								
19:16	O ³	Total Open Resources (TOR): This field indicates the total number of open resources (i.e., allocated and unallocated open resources). If the ZNRM bit is cleared to '0', then this field should be ignored by the host.								
23:20	O ³	Unallocated Open Resources (UOR): This field indicates the number of unallocated open resources. If the ZNRM bit is cleared to '0', then this field should be ignored by the host.								
27:24	O ³	Total ZRWA Resources (TZRWAR): This field indicates the total number of ZRWA resources (i.e., allocated and unallocated ZRWA resources). If the ZNRM bit is cleared to '0', then this field should be ignored by the host.								

Figure 50: I/O Command Set Specific Identify Controller Data Structure for the Zoned Namespace Command Set

Bytes	O/M ¹	Description
31:28	O ³	Unallocated ZRWA Resources (UZRWAR): This field indicates the number of unallocated ZRWA resources. If the ZNRM bit is cleared to '0', then this field should be ignored by the host.
4095:32		Reserved
NOTES: 1. O/M definition: O = Optional, M = Mandatory. 2. Mandatory for controllers that support the Zone Append command. 3. This field shall be supported if the controller supports Zoned Namespace Resource Management (i.e., the ZNRM bit is set to '1').		

4.1.6 Namespace Management command

The Namespace Management command operates as defined in the NVM Command Set Specification with the following changes:

- The addition that the Format Index (refer to section 1.4.2.2) also specifies the entry in the LBA Format Extensions list (refer to Figure 48) for zone attributes used to format a created namespace.
- The Host Software Specified Fields data structure for the Zoned Namespace Command Set is defined in Figure TBDX.
- If the value in the RAR field is greater than the value specified in the UAR field (refer to Figure 50), then the controller shall abort the command with a status code of Invalid Field in Command.
- If the value in the ROR field is greater than the value specified in the UOR field (refer to Figure 50), then the controller shall abort the command with a status code of Invalid Field in Command.
- If the value in the RNUMZRWA field is greater than the value specified in the UZRWAR field (refer to Figure 50), then the controller shall abort the command with a status code of Invalid Field in Command.
- If the values in the MAR field, MOR field, and NUMZRWA field (refer to Figure 48) do not meet the requirements specified in section 2.1.1.4, then the controller shall abort the command with a status code of Invalid Namespace or Format.

After successful creation of a namespace, the resource allocated to the namespace are reflected in the Unallocated Active Resources, Unallocated Open Resources and in the Unallocated ZRWA Resources fields.

Figure TBDX: Namespace Management – Host Software Specified Fields

Bytes	Description	Host Specified
Fields that are a subset of the Identify Namespace data structure		
383:00	Refer to the NVM Command Set Specification	
Fields that are not a subset of the Identify Namespace data structure		
391:384	Refer to the NVM Command Set Specification	Yes
498:392	Reserved	

Figure TBDX: Namespace Management – Host Software Specified Fields

Bytes	Description	Host Specified						
499	Zoned Namespace Create Options	Yes						
	<table><tr><th>Bits</th><th>Description</th></tr><tr><td>07:01</td><td>Reserved</td></tr><tr><td>00</td><td>Allocate ZRWA Resources (AZR): If set to '1', then the namespace is to be created with the number of ZRWA resource specified in the RNUMZRWA field of this data structure. If cleared to '0', then no ZRWA resources are allocated to the namespace to be created. If the ZRWASUP bit is cleared to '0' (refer to Figure 48), then this field shall be ignored by the controller.</td></tr></table>		Bits	Description	07:01	Reserved	00	Allocate ZRWA Resources (AZR): If set to '1', then the namespace is to be created with the number of ZRWA resource specified in the RNUMZRWA field of this data structure. If cleared to '0', then no ZRWA resources are allocated to the namespace to be created. If the ZRWASUP bit is cleared to '0' (refer to Figure 48), then this field shall be ignored by the controller.
	Bits		Description					
	07:01		Reserved					
00	Allocate ZRWA Resources (AZR): If set to '1', then the namespace is to be created with the number of ZRWA resource specified in the RNUMZRWA field of this data structure. If cleared to '0', then no ZRWA resources are allocated to the namespace to be created. If the ZRWASUP bit is cleared to '0' (refer to Figure 48), then this field shall be ignored by the controller.							
503:500	Requested Active Resources (RAR)¹: This field specifies the number of active resources to be allocated to the created namespace. A value of 00000000h specifies the number of Active resources requested to be allocated to the namespace is vendor specific and that number is reflected in the MAR field in the I/O Command Set Specific Identify Namespace data structure. Values of 00000001h to FFFFFFFEh specify the number of Active resources requested to be allocated is as defined in section 2.1.1.4 and that number is reflected in the MAR field in the I/O Command Set Specific Identify Namespace data structure. A value of FFFFFFFFh specifies the number of Active resources requested to be allocated to the namespace is not limited and that number is reflected in the MAR field in the I/O Command Set Specific Identify Namespace data structure.	Yes						
507:504	Requested Open Resources (ROR)¹: This field specifies the number of open resources to be allocated to the created namespace. A value of 00000000h specifies the number of Open resources requested to be allocated to the namespace is vendor specific and that number is reflected in the MOR field in the I/O Command Set Specific Identify Namespace data structure. Values of 00000001h to FFFFFFFEh specify the number of Open resources requested to be allocated is as defined in section 2.1.1.4 and that number is reflected in the MOR field in the I/O Command Set Specific Identify Namespace data structure. A value of FFFFFFFFh specifies the number of Open resources requested to be allocated to the namespace is not limited and that number is reflected in the MOR field in the I/O Command Set Specific Identify Namespace data structure.	Yes						

Figure TBOX: Namespace Management – Host Software Specified Fields

Bytes	Description	Host Specified
511:508	<p>Requested Number of ZRWA Resources (RNUMZRWA)¹: This field specifies the number of ZRWA resources to be allocated to the created namespace.</p> <p>A value of 00000000h specifies the number of ZRWA resources requested to be allocated to the namespace is vendor specific and that number is reflected in the NUMZRWA field in the I/O Command Set Specific Identify Namespace data structure.</p> <p>Values of 00000001h to FFFFFFFEh specify the number of ZRWA resources requested to be allocated is as defined in section 2.1.1.4 and that number is reflected in the NUMZRWA field in the I/O Command Set Specific Identify Namespace data structure.</p> <p>A value of FFFFFFFFh specifies the number of ZRWA resources requested to be allocated to the namespace is not limited and that number is reflected in the NUMZRWA field in the I/O Command Set Specific Identify Namespace data structure.</p>	Yes
<p>Notes:</p> <p>1. If Zoned Namespace Resource Management (refer to the ZNRM bit in Figure 50) is not supported, then this field is ignored by the controller.</p>		

4.1.6.1 Command Completion

When the command is completed, the controller posts a completion queue entry to the Admin Completion Queue indicating the status for the command.

After successful completion of a Namespace Management command with the create operation, the number of resources allocated with the value requested in the RAR field, may be reflected by a change in the UAR field in the Identify Controller data structure by one or more controllers.

After successful completion of a Namespace Management command with the create operation, the number of resources allocated with the value requested in the ROR field, may be reflected by a change in the UOR field in the Identify Controller data structure by one or more controllers.

After successful completion of a Namespace Management command with the create operation, the number of resources allocated with the value requested in the RNUMZRWAS field, may be reflected by a change in the UZRWAR field in the Identify Controller data structure by one or more controllers.