



LEGAL NOTICE:

© Copyright 2007 to 2021 NVM Express™, Inc. ALL RIGHTS RESERVED.

This NVM Express revision 1.4 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 1.4 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 to 2021 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 for New Feature

Technical Proposal ID	4102 ZNS Status Code and Async Notification Enhancement
Change Date	2021-01-14
Builds on Specification	TP 4053a
References Specification	

Technical Proposal Author(s)

Name	Company
Mike Allison, Shirish Bahirat	Intel
Scott Lee, Lee Prewitt	Microsoft

This proposal intends to:

- Allow the Zone Management Send command to return new status codes if the requested new zone state is not achieved due to a controller excursion;
- Allow more than one time for reporting controller recommendations;
- Allow the host to get a list of zones with pending controller recommendations.

Revision History

Revision Date	Change Description
2020-10-30	Initial version
2020-11-03	Per the ZNS Task Group review moving the indexes to its own fields. Editorial edits.
2020-11-09	Fixed Figure 31 Zone if Offline text to use ZSO:Offline versus ZSRO:Read Only.
2020-11-11	Changed “selects” to “indicates”. Changed “zone attribute” to “zone attribute information”.
2020-11-17	Removed the use of “Index”. Updated the summary of the changes.
2020-11-19	Editorial changes in the Description for TP 4053 changes.
2020-11-30	Added a reporting Option note to avoid values used by ZAC/ZBC. Changed ‘;’ to ‘.’ At the end of a list. Decided not to rename the fields as I could not think of better names unless someone has a good suggestion.
2020-12-01	Updated Reporting Options to specifically reserve values used by ZAC/ZBC. Decided that current field names are adequate. Fixed bolding issues. Approved for 30 day member review.
2020-12-03	Accepted all changes and removed all comments.
2021-01-04	Changed date for 2021. Removed section 3.3.1 as no changes were required.
2021-01-05	Changed 2020 ->2021 and corrected the section number from 3 to 3.1. Approved by ZNS Task group for integration request
2021-01-07	Accepted all changes as Technical WG approval for integration.
2021-01-12	Integrated into NVM Express Zoned Namespace Command Set specification.

2021-01-14	Accepted all changes and removed all comments for ratification. Changed 2020 -> 2021.
------------	---

Description for TP 4053 Changes Document

The TP expands TP 4053a by the following changes:

- The Zone Management Send command is allowed to return the status codes Zone Is Offline and Zone Is Read Only to notify the host of excursion that occur during the processing of the command. No asynchronous event occurs when these status codes are utilized.
- The Identify I/O Command Set Specific Namespace Data Structure for the Zoned Namespace Command Set and the Zone Descriptor data structure have been modified to allow a controller to report up to four different time limits for zone recommendations.
- The Zone Receive Action Specific Field in the Zone Management Receive command added a new reporting option for the host to retrieve zones with zone recommendations.

Description of Specification Changes

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.
<Green Bracketed>:	Notes to editor

Modify portions of TP 4053a as shown below:

Modify portions of section 3.1 as shown below:

3.1 Identify Namespace Data Structures

...

Figure 8: Identify I/O Command Set Specific Namespace Data Structure for the Zoned Namespace Command Set

Bytes	O/M ¹	Description								
01:00	O	Zone Operation Characteristics (ZOC): This field indicates the zone operation characteristics of the zoned namespace.								
		<table><tr><th>Bits</th><th>Description</th></tr><tr><td>15:2</td><td>Reserved</td></tr><tr><td>1</td><td>Zone Active Excursions: If set to '1', then a controller may transition a zone in the ZSIO:Implicitly Opened state, the ZSEO:Explicitly Opened state, or the ZSC:Closed state to the ZSF:Full state due to a vendor specific excursion event. If cleared to '0', then a controller shall not transition a zone due to a vendor specific excursion event. Refer to section 5.6.</td></tr><tr><td>0</td><td>Variable Zone Capacity: if set to '1', then the capacity for a zone may change without a change to the format of the zoned namespace. If cleared to '0', then the capacity for a zone does not change without a change to the format of the zoned namespace. Refer to Figure 28.</td></tr></table>	Bits	Description	15:2	Reserved	1	Zone Active Excursions: If set to '1', then a controller may transition a zone in the ZSIO:Implicitly Opened state, the ZSEO:Explicitly Opened state, or the ZSC:Closed state to the ZSF:Full state due to a vendor specific excursion event. If cleared to '0', then a controller shall not transition a zone due to a vendor specific excursion event. Refer to section 5.6.	0	Variable Zone Capacity: if set to '1', then the capacity for a zone may change without a change to the format of the zoned namespace. If cleared to '0', then the capacity for a zone does not change without a change to the format of the zoned namespace. Refer to Figure 28.
		Bits	Description							
		15:2	Reserved							
		1	Zone Active Excursions: If set to '1', then a controller may transition a zone in the ZSIO:Implicitly Opened state, the ZSEO:Explicitly Opened state, or the ZSC:Closed state to the ZSF:Full state due to a vendor specific excursion event. If cleared to '0', then a controller shall not transition a zone due to a vendor specific excursion event. Refer to section 5.6.							
0	Variable Zone Capacity: if set to '1', then the capacity for a zone may change without a change to the format of the zoned namespace. If cleared to '0', then the capacity for a zone does not change without a change to the format of the zoned namespace. Refer to Figure 28.									
03:02	O	Optional Zoned Command Support (OZCS): This field defines optional features of the zoned namespace.								
		<table><tr><th>Bits</th><th>Description</th></tr><tr><td>15:1</td><td>Reserved</td></tr><tr><td>0</td><td>Read Across Zone Boundaries: If set to '1', then any command is allowed to perform read operations that specify an LBA range containing logical blocks in more than one zone. If cleared to '0', then any command that performs a read operation that specifies an LBA range containing logical blocks in more than one zone is aborted as described in section 2.2.2.3.</td></tr></table>	Bits	Description	15:1	Reserved	0	Read Across Zone Boundaries: If set to '1', then any command is allowed to perform read operations that specify an LBA range containing logical blocks in more than one zone. If cleared to '0', then any command that performs a read operation that specifies an LBA range containing logical blocks in more than one zone is aborted as described in section 2.2.2.3.		
		Bits	Description							
15:1	Reserved									
0	Read Across Zone Boundaries: If set to '1', then any command is allowed to perform read operations that specify an LBA range containing logical blocks in more than one zone. If cleared to '0', then any command that performs a read operation that specifies an LBA range containing logical blocks in more than one zone is aborted as described in section 2.2.2.3.									
07:04	M	Maximum Active Resources (MAR): This field defines the maximum number of zones that are allowed to be in the ZSIO:Implicitly Opened state, the ZSEO:Explicitly Opened state, or the ZSC:Closed state. A value of FFFFFFFh indicates that there is no limit. This is a 0's based value								
11:08	M	Maximum Open Resources (MOR): This field defines the maximum number of zones that are allowed to be in the ZSIO:Implicitly Opened state or the ZSEO:Explicitly Opened state. The Maximum Open Resource field shall be less than or equal to the Maximum Active Resources field. A value of FFFFFFFFh indicates that there is no limit. This is a 0's based value								
15:12	O	Reset Recommended Limit (RRL): If the zone attribute Reset Zone Recommended Time Limit field is set to 00b, then this field indicates the number of seconds before the NVM subsystem may perform a vendor specific action on a zone after the Reset Zone Recommended zone attribute is set to '1' for that zone. If this field is cleared to 0h, then no Reset Recommended Limit is reported. Refer to section 5.4.								
19:16	O	Finish Recommended Limit (FRL): If the zone attribute Finish Zone Recommended Time Limit field is set to 00b, then this field indicates the number of seconds before the NVM subsystem may perform the vendor specific action on a zone after the Finish Zone Recommended zone attribute is set to '1' for that zone. If this field is cleared to 0h, then no Finish Recommended Limit is reported. Refer to section 5.5.								
23:20	O	Reset Recommended Limit 1 (RRL 1): If the zone attribute Reset Zone Recommended Time Limit field is set to 01b, then this field indicates the number of seconds before the NVM subsystem may perform a vendor specific action on a zone after the Reset Zone Recommended zone attribute is set to '1' for that zone. If this field is cleared to 0h, then no Reset Recommended Limit is reported. Refer to section 5.4.								

Bytes	O/M ¹	Description
27:24	O	Reset Recommended Limit 2 (RRL 2): If the zone attribute Reset Zone Recommended Time Limit field is set to 10b, then this field indicates the number of seconds before the NVM subsystem may perform a vendor specific action on a zone after the Reset Zone Recommended zone attribute is set to '1' for that zone. If this field is cleared to 0h, then no Reset Recommended Limit is reported. Refer to section 5.4.
31:28	O	Reset Recommended Limit 3 (RRL 3): If the zone attribute Reset Zone Recommended Time Limit field is set to 11b, then this field indicates the number of seconds before the NVM subsystem may perform a vendor specific action on a zone after the Reset Zone Recommended zone attribute is set to '1' for that zone. If this field is cleared to 0h, then no Reset Recommended Limit is reported. Refer to section 5.4.
35:32	O	Finish Recommended Limit 1 (FRL1): If the zone attribute Finish Zone Recommended Time Limit field is set to 01b, then this field indicates the number of seconds before the NVM subsystem may perform the vendor specific action on a zone after the Finish Zone Recommended zone attribute is set to '1' for that zone. If this field is cleared to 0h, then no Finish Recommended Limit is reported. Refer to section 5.5.
39:36	O	Finish Recommended Limit 2 (FRL2): If the zone attribute Finish Zone Recommended Time Limit field is set to 10b, then this field indicates the number of seconds before the NVM subsystem may perform the vendor specific action on a zone after the Finish Zone Recommended zone attribute is set to '1' for that zone. If this field is cleared to 0h, then no Finish Recommended Limit is reported. Refer to section 5.5.
43:40	O	Finish Recommended Limit 3 (FRL3): If the zone attribute Finish Zone Recommended Time Limit field is set to 11b, then this field indicates the number of seconds before the NVM subsystem may perform the vendor specific action on a zone after the Finish Zone Recommended zone attribute is set to '1' for that zone. If this field is cleared to 0h, then no Finish Recommended Limit is reported. Refer to section 5.5.
2815:4420		Reserved
2831:2816	M	LBA Format 0 Extension (LBAFE0): This field indicates the LBA format Extension 0 that is supported by the controller. The Zone format field is defined in Figure 47.
2847:2832	O	LBA Format 1 Extension (LBAFE1): This field indicates the LBA format Extension 1 that is supported by the controller. The LBA Format Extension field is defined in Figure 47.
...		
3823:3808	O	LBA Format 62 Extension (LBAFE62): This field indicates the LBA format Extension 62 that is supported by the controller. The LBA Format Extension field is defined in Figure 47.
3839:3824	O	LBA Format 63 Extension (LBAFE63): This field indicates the LBA format Extension 63 that is supported by the controller. The LBA Format Extension field is defined in Figure 47.
4095:3840	O	Vendor Specific
NOTES:		
1. O/M definition: O = Optional, M = Mandatory.		

Modify portions of section 4.3.2 as shown below:

4.3.2 Command Completion

When the command is completed, the controller shall post a completion queue entry to the associated I/O Completion Queue indicating the status for the command. Figure 31 defines the Zone Management Send command specific status values.

Figure 31: Zone Management Send – Command Specific Status Values

Value	Description
BAh	Zone Is Read Only: Zone is in the ZSRO:Read Only state. This may have occurred during the processing of the command.
BBh	Zone Is Offline: Zone is in the ZSO:Offline state. This may have occurred during the processing of the command.
BFh	Invalid Zone State Transition: The request is not a valid zone state transition.

Dword 0 of the completion queue entry indicates if the zone capacity of the specified zone has been modified. The definition of Dword 0 of the completion queue entry is in Figure 32.

Figure 32: Zone Management Send – Completion Queue Entry Dword 0

Bits	Description
31:01	Reserved
00	Zone Capacity Changed: This bit indicates if the zone capacity has changed. If this bit is set to '1', then the zone capacity has changed due to this command. The host should read the Zone Descriptor data structure for the zone specified by the SLBA field. If this bit is cleared to '0', then the zone capacity has not changed due to this command.

Modify portions of Figure 21 in section 4.1 as shown below:

4.1 Command Specific Status Codes

...

- Figure 21: Status Code – Command Specific Status Values, Zoned Namespace Command Set**

Value	Description	Commands Affected
B8h	Zone Boundary Error	Read ¹ , Compare ¹ , Verify ¹ , Copy ¹ , Write, Write Uncorrectable, Write Zeroes, Zone Append
B9h	Zone Is Full	Write, Write Uncorrectable, Write Zeroes, Copy, Zone Append
BAh	Zone Is Read Only	Write, Write Uncorrectable, Write Zeroes, Copy, Zone Append, Zone Management Send
BBh	Zone Is Offline	Read, Compare, Verify, Copy, Write, Write Uncorrectable, Write Zeroes, Zone Append, Zone Management Send
BCh	Zone Invalid Write	Write, Write Uncorrectable, Write Zeroes, Copy
BDh	Too Many Active Zones	Write, Write Uncorrectable, Write Zeroes, Copy, Zone Append, Zone Management Send
BEh	Too Many Open Zones	Write, Write Uncorrectable, Write Zeroes, Copy, Zone Append, Zone Management Send
BFh	Invalid Zone State Transition	Zone Management Send
NOTES:		
1. This command is affected if the Read Across Zone Boundaries bit in the Zoned Namespace Command Set specific Identify Namespace data structure is cleared to '0'.		

Modify portions of figure 36 in section 4.4 as shown below:

4.4 Zone Management Receive command

...

Bits	Description
	Value Description
	00h Report Zones: Reports zone descriptor entries through the Report Zones data structure (refer to Figure 37).
	01h Extended Report Zones: Reports zone descriptor entries through the Extended Report Zones data structure (refer to Figure 38). This value is supported if the zoned namespace is formatted with a non-zero Zone Descriptor Extension Size. Otherwise, the command shall be aborted with a status code of Invalid Field in Command.
	02h to FFh Reserved

Modify portions of section 4.4.2.3 as shown below:

4.4.2.3 Zone Descriptor Data Structure

Figure 39 defines the Zone Descriptor data structure.

Figure 39: Zone Descriptor Data Structure

Bytes	Description																				
00	Bits	Description																			
	7:4	Reserved.																			
	3:0	Zone Type (ZT): This field indicates the type of the zone.																			
		<table><tr><th>Value</th><th>Definition</th></tr><tr><td>2h</td><td>Sequential Write Required</td></tr><tr><td>All other values</td><td>Reserved</td></tr></table>	Value	Definition	2h	Sequential Write Required	All other values	Reserved													
Value		Definition																			
2h	Sequential Write Required																				
All other values	Reserved																				
01	Bits	Description																			
	7:4	Zone State (ZS): This field indicates the state of the zone.																			
		<table><tr><th>Value</th><th>Definition</th></tr><tr><td>0h</td><td>Reserved</td></tr><tr><td>1h</td><td>Empty</td></tr><tr><td>2h</td><td>Implicitly Opened</td></tr><tr><td>3h</td><td>Explicitly Opened</td></tr><tr><td>4h</td><td>Closed</td></tr><tr><td>5h to Ch</td><td>Reserved</td></tr><tr><td>Dh</td><td>Read Only</td></tr><tr><td>Eh</td><td>Full</td></tr><tr><td>Fh</td><td>Offline</td></tr></table>	Value	Definition	0h	Reserved	1h	Empty	2h	Implicitly Opened	3h	Explicitly Opened	4h	Closed	5h to Ch	Reserved	Dh	Read Only	Eh	Full	Fh
Value		Definition																			
0h		Reserved																			
1h		Empty																			
2h		Implicitly Opened																			
3h		Explicitly Opened																			
4h		Closed																			
5h to Ch		Reserved																			
Dh		Read Only																			
Eh	Full																				
Fh	Offline																				
3:0	Reserved																				
02	Zone Attributes (ZA): Indicates attributes for the Zone:																				
	Bits	Description																			
	7	Zone Descriptor Extension Valid (ZDEV): If this bit is set to '1', then Zone Descriptor Extension data is associated with the zone. If this bit is cleared to '0', then no Zone Descriptor Extension data is associated with the zone. Refer to section 5.1.																			
	6:3	Reserved																			
	2	Reset Zone Recommended (RZR): If this bit is set to '1', then the controller recommends that this zone be reset. Refer to section 5.2.																			
1	Finish Zone Recommended (FZR): If this bit is set to '1', then the controller recommends that this zone be finished. Refer to section 5.3.																				

Bytes	Description								
0	Zone Finished by Controller (ZFC): If this bit is set to '1', then the controller finished this zone due to a Zone Active Excursion. Refer to section 5.4.								
03	Zone Attributes Information (ZAI): Indicates additional information associated with attributes for the Zone: <table> <tr> <th>Bits</th><th>Description</th></tr> <tr> <td>7:4</td><td>Reserved</td></tr> <tr> <td>3:2</td><td> Reset Zone Recommended Time Limit (RZRTL): If the Reset Zone Recommended bit is set to '1', then the value in this field selects a field in the Identify I/O Command Set Specific namespace data structure for the Zoned Namespace Command Set that indicates amount of time before the NVM subsystem may perform a vendor specific action on a zone after the Reset Zone Recommended zone attribute is set to '1' for that zone: <ul style="list-style-type: none"> A 00b value indicates the Reset Recommended Limit (RRL) field; A 01b value indicates the Reset Recommended Limit 1 (RRL1) field; A 10b value indicates the Reset Recommended Limit 2 (RRL2) field; and A 11b value indicates the Reset Recommended Limit 3 (RRL3) field. </td></tr> <tr> <td>1:0</td><td> Finish Zone Recommended Time Limit (FZRTL): If the Finish Zone Recommended bit is set to '1', then the value in this field selects a field in the Identify I/O Command Set Specific namespace data structure for the Zoned Namespace Command Set that indicates the amount of time before the NVM subsystem may perform a vendor specific action on a zone after the Finish Zone Recommended zone attribute is set to '1' for that zone: <ul style="list-style-type: none"> A 00b value indicates the Finish Recommended Limit (FRL) field; A 01b value indicates the Finish Recommended Limit 1 (FRL1) field; A 10b value indicates the Finish Recommended Limit 2 (FRL2) field; and A 11b value indicates the Finish Recommended Limit 3 (FRL3) field. </td></tr> </table>	Bits	Description	7:4	Reserved	3:2	Reset Zone Recommended Time Limit (RZRTL): If the Reset Zone Recommended bit is set to '1', then the value in this field selects a field in the Identify I/O Command Set Specific namespace data structure for the Zoned Namespace Command Set that indicates amount of time before the NVM subsystem may perform a vendor specific action on a zone after the Reset Zone Recommended zone attribute is set to '1' for that zone: <ul style="list-style-type: none"> A 00b value indicates the Reset Recommended Limit (RRL) field; A 01b value indicates the Reset Recommended Limit 1 (RRL1) field; A 10b value indicates the Reset Recommended Limit 2 (RRL2) field; and A 11b value indicates the Reset Recommended Limit 3 (RRL3) field. 	1:0	Finish Zone Recommended Time Limit (FZRTL): If the Finish Zone Recommended bit is set to '1', then the value in this field selects a field in the Identify I/O Command Set Specific namespace data structure for the Zoned Namespace Command Set that indicates the amount of time before the NVM subsystem may perform a vendor specific action on a zone after the Finish Zone Recommended zone attribute is set to '1' for that zone: <ul style="list-style-type: none"> A 00b value indicates the Finish Recommended Limit (FRL) field; A 01b value indicates the Finish Recommended Limit 1 (FRL1) field; A 10b value indicates the Finish Recommended Limit 2 (FRL2) field; and A 11b value indicates the Finish Recommended Limit 3 (FRL3) field.
Bits	Description								
7:4	Reserved								
3:2	Reset Zone Recommended Time Limit (RZRTL): If the Reset Zone Recommended bit is set to '1', then the value in this field selects a field in the Identify I/O Command Set Specific namespace data structure for the Zoned Namespace Command Set that indicates amount of time before the NVM subsystem may perform a vendor specific action on a zone after the Reset Zone Recommended zone attribute is set to '1' for that zone: <ul style="list-style-type: none"> A 00b value indicates the Reset Recommended Limit (RRL) field; A 01b value indicates the Reset Recommended Limit 1 (RRL1) field; A 10b value indicates the Reset Recommended Limit 2 (RRL2) field; and A 11b value indicates the Reset Recommended Limit 3 (RRL3) field. 								
1:0	Finish Zone Recommended Time Limit (FZRTL): If the Finish Zone Recommended bit is set to '1', then the value in this field selects a field in the Identify I/O Command Set Specific namespace data structure for the Zoned Namespace Command Set that indicates the amount of time before the NVM subsystem may perform a vendor specific action on a zone after the Finish Zone Recommended zone attribute is set to '1' for that zone: <ul style="list-style-type: none"> A 00b value indicates the Finish Recommended Limit (FRL) field; A 01b value indicates the Finish Recommended Limit 1 (FRL1) field; A 10b value indicates the Finish Recommended Limit 2 (FRL2) field; and A 11b value indicates the Finish Recommended Limit 3 (FRL3) field. 								
07:0304	Reserved								
15:08	Zone Capacity (ZCAP): This field contains the maximum number of logical blocks that are available to be written with user data when the zone is in the ZSE:Empty state. This value shall be less than or equal to the Zone Size field (refer to Figure 9). If the Variable Zone Capacity bit is cleared to '0' in the Zone Operation Characteristics field in the Zoned Namespace Command Set specific Identify Namespace data structure (refer to Figure 8), then this field does not change without a change to the format of the zoned namespace. If the Variable Zone Capacity bit is set to '1' in the Zone Operation Characteristics field in the Zoned Namespace Command Set specific Identify Namespace data structure (refer to Figure 8), then the zone capacity may change upon successful completion of a Zone Management Send command specifying the Zone Send Action of Reset Zone (refer to section 4.3.1.4).								
23:16	Zone Start Logical Block Address (ZSLBA): This field contains the 64-bit address of the lowest logical block for a zone.								
31:24	Write Pointer (WP): This field is the logical block address where the next write operation for this zone should be issued. Refer to section 2.3.1 for the behavior of the write pointer.								
63:32	Reserved								

Modify portions of section 5.2 as shown below:

5.2 Reset Zone Recommended

A controller that schedules an internal operation (e.g., background operation on the non-volatile media) on a zone that is in the ZSF:Full state may notify host software to perform a zone reset operation (refer to section 4.3.1.4). If a controller schedules such an internal operation on a zone, the controller may notify the host by:

- setting the Reset Zone Recommended zone attribute of the specific zone to '1' (refer to Figure 39);

- b) setting the Reset Zone Recommended Time Limit zone attribute information to indicate the number of seconds before the controller intends to perform an internal operation on the specified zone; and
- c) generating a Zone Descriptor Changed event for the specific zone.

~~If the Reset Zone Recommended zone attribute of a specific zone is set to '1', the Reset Recommended Limit field (refer to Figure 10) indicates the number of seconds before the controller intends to perform an internal operation on the specified zone.~~

As a zone reset operation destroys data in a specific zone, it is optional for the host software to perform a zone reset operation on zones that have the Reset Zone Recommended zone attribute set to '1'. If the host does not perform a zone reset operation on the specific zone, then the internal operation, which may impact performance, may be performed.

If the controller has processed the internal operation or the internal operation is no longer scheduled, the controller may notify the host by:

- a) clearing the Reset Zone Recommended zone attribute of the specific zone to '0'; and
- b) generating a Zone Descriptor Changed event for the specific zone

Modify portions of section 5.3 as shown below:

5.3 Finish Zone Recommended

A controller that schedules an internal operation (e.g., background operation on the non-volatile media) on a zone that is in the:

- a) ZSIO:Implicitly Opened state;
- b) ZSEO:Explicitly Opened state; or
- c) ZSC:Closed state

may notify host software to either:

- a) perform a zone finish operation (reset to section 4.3.1.2) on that zone; or
- b) initiate write operations to that zone such that the zone transitions to the ZSF:Full state,

so that the controller may cancel the scheduled internal operation. If a controller schedules such an internal operation on a zone, the controller may notify the host by:

- a) setting the Finish Zone Recommended zone attribute of the specific zone to '1' (refer to Figure 39);
- b) setting the Finish Zone Recommended Time Limit zone attribute information to indicate the number of seconds before the controller intends to perform an internal operation on the specified zone; and
- c) generating a Zone Descriptor Changed event for the specific zone.

~~If the Finish Zone Recommended zone attribute of a specific zone is set to '1', the Finish Recommended Limit field (refer to Figure 10) indicates the number of seconds before the controller intends to perform an internal operation on the specified zone.~~

It is optional for the host software to process the above mitigating actions on zones that have the Finish Zone Recommended zone attribute set to '1'. If the host does not perform the mitigating actions on the specific zone, then the internal operation, which may impact performance, may be performed.

If the controller has processed the internal operation or the internal operation is no longer scheduled, the controller may notify the host by:

- a) clearing the Finish Zone Recommended zone attribute of the specific zone to '0'; and
- b) generating a Zone Descriptor Changed event for the specific zone.