

---

## **Information technology - Zoned Device ATA Command Set (ZAC)**

This is a draft proposed American National Standard of Accredited Standards Committee INCITS. As such this is not a completed standard. The T13 Technical Committee may modify this document as a result of comments received during public review and its approval as a standard. Use of the information contained here in is at your own risk.

Permission is granted to members of INCITS, its technical committees, and their associated task groups to reproduce this document for the purposes of INCITS standardization activities without further permission, provided this notice is included. All other rights are reserved. Any commercial or for-profit replication or republication is prohibited.

T13 Technical Editor:

Ralph O. Weber  
Western Digital Technologies, Inc.  
18484 Preston Road, Suite 102, PMB 178  
Dallas, TX 75252  
USA

Telephone: 214-912-1373  
Email: Ralph.Weber at wdc.com

## Points of Contact

### T13 Chair

Dan Colgrove  
Toshiba America Electronic Components, Inc.  
2590 Orchard Parkway  
San Jose, CA 95131  
Tel: 408-526-2563

### T13 Vice-Chair

Jim Hatfield  
Seagate Technology  
389 Disc Drive  
Longmont CO 80503  
Tel: 720-684-2120

### INCITS Secretariat

INCITS Secretariat  
1101 K Street NW Suite 610  
Washington, DC 20005  
Email: [INCITS@ITIC.ORG](mailto:INCITS@ITIC.ORG)  
Tel: 202-737-8888  
Fax: 202-638-4922

### T13 Reflector

See the T13 Web Site at <http://www.t13.org> for reflector information.

### T13 Web Site

<http://www.t13.org>

### T13 FTP Site

<ftp.t13.org> (see [www.t13.org](http://www.t13.org) for login information)

### Purchase INCITS Standards

<http://www.incits.org/standards-information/purchase-standards-or-download-dpans>

# Draft

Secretariat  
Information Technology Industry Council

Approved mm.dd.yy  
American National Standards Institute, Inc.

## **ABSTRACT**

This standard specifies the command set that host systems use to access storage devices that implement one of the zones feature sets. This provides a common command set for systems manufacturers, system integrators, software suppliers, and suppliers of storage devices that provide one of the zones feature sets.

## American National Standard

Approval of an American National Standard requires verification by ANSI that the requirements for due process, consensus, and other criteria for approval have been met by the standards developer. Consensus is established when, in the judgment of the ANSI Board of Standards Review, substantial agreement has been reached by directly and materially affected interests. Substantial agreement means much more than a simple majority, but not necessarily unanimity. Consensus requires that all views and objections be considered, and that effort be made towards their resolution.

The use of American National Standards is completely voluntary; their existence does not in any respect preclude anyone, whether he has approved the standards or not, from manufacturing, marketing, purchasing, or using products, processes, or procedures not conforming to the standards.

The American National Standards Institute does not develop standards and will in no circumstances give interpretation on any American National Standard. Moreover, no person shall have the right or authority to issue an interpretation of an American National Standard in the name of the American National Standards Institute. Requests for interpretations should be addressed to the secretariat or sponsor whose name appears on the title page of this standard.

**CAUTION NOTICE:** This American National Standard may be revised or withdrawn at any time. The procedures of the American National Standards Institute require that action be taken periodically to reaffirm, revise, or withdraw this standard. Purchasers of American National Standards may receive current information on all standards by calling or writing the American National Standards Institute.

CAUTION: The developers of this standard have requested that holders of patents that may be required for the implementation of the standard, disclose such patents to the publisher. However, neither the developers nor the publisher have undertaken a patent search in order to identify which, if any, patents may apply to this standard. As of the date of publication of this standard, following calls for the identification of patents that may be required for the implementation of the standard, no such claims have been made. No further patent search is conducted by the developer or the publisher in respect to any standard it processes. No representation is made or implied that licenses are not required to avoid infringement in the use of this standard.

Published by

**American National Standards Institute**  
**23 W. 43rd Street, New York, New York 10036**

Copyright © 2015 by Information Technology Industry Council (ITI).  
All rights reserved.

No part of this publication may be reproduced in any form, in an electronic retrieval system or otherwise, without prior written permission of:

Information Technology Industry Council  
1101 K St, NW Suite 610  
Washington, D.C. 20005

Printed in the United States of America

## Contents

	Page
Points of Contact.....	ii
Contents.....	v
Tables.....	ix
Figures.....	x
Foreword.....	xi
Introduction.....	xiv
1 Scope.....	1
2 Normative references.....	2
3 Definitions, abbreviations, and conventions.....	3
3.1 Definitions.....	3
3.2 Symbols and abbreviations.....	7
3.2.1 Abbreviations.....	7
3.2.2 Units.....	8
3.2.3 Symbols.....	8
3.2.4 Mathematical operators.....	8
3.3 Keywords.....	8
3.4 Conventions.....	9
3.4.1 Overview.....	9
3.4.2 Precedence.....	10
3.4.3 Lists.....	10
3.4.3.1 Lists overview.....	10
3.4.3.2 Unordered lists.....	10
3.4.3.3 Ordered lists.....	10
3.4.4 Numbering.....	11
3.4.5 Bit conventions.....	11
3.4.6 Number range convention.....	11
3.4.7 State diagram convention.....	12
3.4.8 Offset convention.....	12
4 Feature set definitions.....	13
4.1 Introduction.....	13
4.2 Basic zoned device model.....	13
4.3 Host Aware Zones feature set.....	13
4.4 Host Managed Zones feature set.....	14
4.5 Zone attributes.....	15
4.5.1 Summary.....	15
4.5.2 Zone Type zone attribute.....	15
4.5.3 Zone Condition zone attribute.....	16
4.5.4 WPointer zone attribute.....	16
4.5.5 RWP Recommended zone attribute.....	16
4.5.6 Non-Sequential Write Resources Active zone attribute.....	16
4.6 Zone types.....	17
4.6.1 Overview.....	17
4.6.2 Conventional zones.....	17
4.6.3 Write pointer zones.....	17
4.6.3.1 Features common to all write pointer zones.....	17

4.6.3.1.1	General .....	17
4.6.3.1.2	Resetting the write pointer .....	19
4.6.3.2	Sequential Write Preferred zones .....	19
4.6.3.2.1	Writing in Sequential Write Preferred zones .....	19
4.6.3.2.2	Reading in Sequential Write Preferred zones .....	21
4.6.3.2.3	Substitute data pattern for Sequential Write Preferred zones .....	21
4.6.3.3	Sequential Write Required zones .....	21
4.6.3.3.1	Writing in Sequential Write Required zones .....	21
4.6.3.3.2	Reading in Sequential Write Required zones .....	22
4.6.3.3.3	Substitute data pattern for Sequential Write Required zones .....	23
4.6.3.3.4	Opening Sequential Write Required zones .....	24
4.6.3.4	Zone Condition state machine .....	24
4.6.3.4.1	Overview .....	24
4.6.3.4.2	ZC1: Empty state .....	27
4.6.3.4.3	ZC2: Implicit_Open state .....	28
4.6.3.4.4	ZC3: Explicit_Open state .....	29
4.6.3.4.5	ZC4: Closed state .....	30
4.6.3.4.6	ZC5: Full state .....	31
4.6.3.4.7	ZC6: Read_Only state .....	31
4.6.3.4.8	ZC7: Offline state .....	32
4.6.3.4.9	Manage Open Zone Resources function .....	32
4.6.3.4.9.1	Overview .....	32
4.6.3.4.9.2	Processing a sequential write preferred zone .....	32
4.6.3.4.9.3	Processing a sequential write required zone .....	33
4.6.3.4.10	Explicitly Open Zone function .....	33
4.6.3.4.11	Close Zone function .....	34
4.6.3.4.12	Finish Zone function .....	34
4.6.3.4.13	Reset Write Pointer function .....	34
4.7	Additional features for the Host Aware Zones feature set and the Host Managed Zones feature set	34
4.7.1	Introduction .....	34
4.7.2	Zoned device internal resource management .....	35
4.7.3	Unexpected power removal .....	35
4.7.4	Media failure .....	36
4.7.5	Interactions with volatile caches .....	36
4.7.6	Interactions with the Sanitize Device feature set .....	37
4.7.7	Interactions with the Power Management feature set .....	38
4.7.8	Interactions with the SECURITY ERASE UNIT command .....	38
4.7.9	Interactions with the SCT Write Same command .....	38
4.7.9.1	Overview .....	38
4.7.9.2	WRITE AND SET SUBSTITUTE DATA PATTERN bit .....	38
4.7.9.3	ZONE POST PROCESSING field .....	39
5	Command descriptions .....	40
5.1	Command description introduction .....	40
5.2	Zone management commands .....	43
5.2.1	Overview .....	43
5.2.2	ZAC Management In command template, 4Ah, DMA .....	44
5.2.2.1	Feature Set .....	44
5.2.2.2	Description .....	44
5.2.2.3	Inputs .....	44
5.2.2.3.1	Overview .....	44
5.2.2.3.2	Action specific fields .....	44
5.2.2.3.3	ZM_ACTION field .....	45
5.2.2.3.4	RETURN PAGE COUNT field .....	45
5.2.2.4	Normal Outputs and Error Outputs .....	45
5.2.2.5	NCQ encapsulation for a DMA ZAC Management In command .....	45
5.2.3	ZAC Management Out command template, 9Fh, Non-Data or DMA .....	46

5.2.3.1 Feature Set .....	46
5.2.3.2 Description .....	46
5.2.3.3 Inputs .....	46
5.2.3.3.1 Overview .....	46
5.2.3.3.2 PAGES TO SEND field .....	46
5.2.3.4 Normal Outputs and Error Outputs .....	47
5.2.3.5 Non-Data NCQ encapsulation .....	47
5.2.3.6 NCQ encapsulation for a DMA ZAC Management Out command .....	47
5.2.4 CLOSE ZONE EXT command – 9Fh/01h, Non-Data .....	48
5.2.4.1 Feature Set .....	48
5.2.4.2 Description .....	48
5.2.4.3 Inputs .....	48
5.2.4.3.1 Overview .....	48
5.2.4.3.2 CLOSE ALL bit .....	48
5.2.4.3.3 ZONE ID field .....	49
5.2.4.4 Normal Outputs .....	49
5.2.4.5 Error Outputs .....	49
5.2.5 FINISH ZONE EXT command – 9Fh/02h, Non-Data .....	50
5.2.5.1 Feature Set .....	50
5.2.5.2 Description .....	50
5.2.5.3 Inputs .....	50
5.2.5.3.1 Overview .....	50
5.2.5.3.2 FINISH ALL bit .....	50
5.2.5.4 Normal Outputs .....	51
5.2.5.5 Error Outputs .....	51
5.2.6 OPEN ZONE EXT command – 9Fh/03h, Non-Data .....	52
5.2.6.1 Feature Set .....	52
5.2.6.2 Description .....	52
5.2.6.3 Inputs .....	52
5.2.6.3.1 Overview .....	52
5.2.6.3.2 OPEN ALL bit .....	52
5.2.6.4 Normal Outputs .....	53
5.2.6.5 Error Outputs .....	53
5.2.7 REPORT ZONES EXT command – 4Ah/00h, DMA .....	54
5.2.7.1 Feature Set .....	54
5.2.7.2 Description .....	54
5.2.7.3 Inputs .....	54
5.2.7.3.1 Overview .....	54
5.2.7.3.2 PARTIAL bit .....	54
5.2.7.3.3 REPORTING OPTIONS field .....	55
5.2.7.3.4 ZONE LOCATOR field .....	55
5.2.7.4 Normal Outputs .....	55
5.2.7.5 Error Outputs .....	55
5.2.7.6 Input From the Device to the Host Data Structure .....	56
5.2.7.6.1 Overview .....	56
5.2.7.6.2 ZONE LIST LENGTH field .....	56
5.2.7.6.3 SAME field .....	57
5.2.7.6.4 MAXIMUM LBA field .....	57
5.2.7.6.5 Zero padding .....	57
5.2.7.6.6 Zone descriptor format .....	58
5.2.7.6.6.1 Overview .....	58
5.2.7.6.6.2 ZONE TYPE field .....	58
5.2.7.6.6.3 ZONE CONDITION field .....	59
5.2.7.6.6.4 NON_SEQ bit .....	59
5.2.7.6.6.5 RESET bit .....	59
5.2.7.6.6.6 ZONE LENGTH field .....	59
5.2.7.6.6.7 ZONE START LBA field .....	59

5.2.7.6.6.8 WRITE POINTER LBA field.....	59
5.2.8 RESET WRITE POINTER EXT command – 9Fh/04h, Non-Data .....	60
5.2.8.1 Feature Set.....	60
5.2.8.2 Description.....	60
5.2.8.3 Inputs.....	60
5.2.8.3.1 Overview .....	60
5.2.8.3.2 RESET ALL bit.....	60
5.2.8.4 Normal Outputs.....	61
5.2.8.5 Error Outputs .....	61
6 Log definitions .....	62
6.1 Introduction .....	62
6.2 IDENTIFY DEVICE data log (Log Address 30h) .....	62
6.2.1 Overview .....	62
6.2.2 Zoned Device Information (page 09h) .....	62
6.2.2.1 Overview.....	62
6.2.2.2 URSWRZ bit .....	63
6.2.2.3 OPTIMAL NUMBER OF OPEN SEQUENTIAL WRITE PREFERRED ZONES field.....	63
6.2.2.4 OPTIMAL NUMBER OF NON-SEQUENTIALLY WRITTEN SEQUENTIAL WRITE PREFERRED ZONES field	63
6.2.2.5 MAXIMUM NUMBER OF OPEN SEQUENTIAL WRITE REQUIRED ZONES field .....	63
6.2.2.6 Version information.....	64
6.2.2.6.1 ZAC MINOR VERSION field .....	64
7 Normal and Error Outputs .....	65
7.1 Overview .....	65
7.2 Normal Outputs .....	65
7.3 Error Outputs .....	66
Annex A (informative) Host Considerations for Zoned Devices .....	67
A.1 Introduction .....	67
A.2 Writing to write pointer zones .....	67
A.3 Open zone considerations .....	67
A.3.1 Overview .....	67
A.3.2 Explicitly opened zones and implicitly opened zones .....	68
A.3.3 Opening and closing zones .....	69
A.3.4 Finish Zone function considerations .....	69
A.4 Open zone resources considerations based on zone type .....	69
A.4.1 Sequential Write Preferred zones .....	69
A.4.2 Sequential Write Required zones .....	70
A.5 Partial failures .....	70
A.5.1 Overview .....	70
A.5.2 Sanitize considerations .....	70
Bibliography .....	72



**Tables**

	Page
Table 1 - Numbering conventions .....	11
Table 2 - Summary of zone attributes .....	15
Table 3 - Zone Type zone attribute .....	15
Table 4 - Relationships between zone attributes .....	15
Table 5 - Zone Condition zone attribute .....	16
Table 6 - Zone Condition state machine functions .....	25
Table 7 - Characteristics and attributes associated with zone state .....	25
Table 8 - ZAC OPTIONS field .....	38
Table 9 - ZONE POST PROCESSING field .....	39
Table 10 - Example Command Structure .....	40
Table 11 - Example Normal Output .....	41
Table 12 - Example Error Output .....	42
Table 13 - Zone management commands .....	43
Table 14 - ZAC Management In command template inputs .....	44
Table 15 - ZM_ACTION field .....	45
Table 16 - RECEIVE FPDMA QUEUED command encapsulation for the inputs used by the ZAC Management In command .....	45
Table 17 - ZAC Management Out command template inputs .....	46
Table 18 - NCQ NON-DATA command encapsulation for the inputs used by the ZAC Management Out command .....	47
Table 19 - SEND FPDMA QUEUED command encapsulation for the inputs used by the ZAC Management Out command .....	47
Table 20 - CLOSE ZONE EXT command inputs .....	48
Table 21 - FINISH ZONE EXT command inputs .....	50
Table 22 - OPEN ZONE EXT command inputs .....	52
Table 23 - REPORT ZONES EXT command inputs .....	54
Table 24 - REPORTING OPTIONS field .....	55
Table 25 - REPORT ZONES EXT input from device to host .....	56
Table 26 - SAME field .....	57
Table 27 - Zone descriptor format .....	58
Table 28 - ZONE TYPE field .....	58
Table 29 - ZONE CONDITION field .....	59
Table 30 - RESET WRITE POINTER EXT command inputs .....	60
Table 31 - Zoned Device Information .....	62
Table 32 - ZAC MINOR VERSION field .....	64
Table 33 - Generic ZAC Extended Normal Output .....	65
Table 34 - Generic ZAC Extended Error Output .....	66

## Figures

	Page
Figure 1 - ATA document relationships .....	1
Figure 2 - State diagram convention .....	12
Figure 3 - Zones in a zoned device .....	13
Figure 4 - Write pointer zone and write pointer after Reset Write Pointer function with no subsequent writes .....	17
Figure 5 - Write pointer zone and write pointer example .....	18
Figure 6 - Example effects of host actions on a write pointer zone .....	19
Figure 7 - Example write command that starts at the write pointer .....	20
Figure 8 - Examples of write commands that do not start at the write pointer .....	20
Figure 9 - Zone Condition state machine .....	26

**Foreword**

(This foreword is not part of this standard.)

This standard specifies the command set that host systems use to access storage devices that implement one of the zones feature sets. This provides a common command set for systems manufacturers, system integrators, software suppliers, and suppliers of storage devices that provide one of the zones feature sets.

Requests for interpretation, suggestions for improvement and addenda, or defect reports are welcome. They should be sent to the INCITS Secretariat, ITI, 1101 K Street NW, Suite 610, Washington, DC 20005.

This standard was processed and approved for submittal to ANSI by InterNational Committee for Information Technology Standards (INCITS). Committee approval of this standard does not necessarily imply that all committee members voted for approval. At the time it approved this standard, INCITS had the following members:

<b>Organization</b>
AIM Global Inc
Adobe Systems Inc
American National Standards Institute
Apple
Department of Commerce - NIST
Distributed Management Task Force (DMTF)
EMC Corporation
Farance Inc
Futurewei Technologies Inc
GS1GO
Hewlett-Packard Enterprise
Hewlett-Packard, Inc.
IBM Corporation
IEEE
Intel Corporation
Microsoft Corporation
Oracle
Purdue University
Telecommunications Industry Association (TIA)
United States Dept of Homeland Security

Technical Committee T13 on ATA Interfaces, that reviewed this standard, had the following members and additional participants:

Dan Colegrove (Toshiba America Electronic Components), Chair

Jim Hatfield (Seagate), Vice-Chair

William Martin (Samsung), Secretary

Company	Name
Avago	George Penokie
	Brad Besmer
Dell	Kevin Marks
	Frank Widjaja
	Mike Wilson
EMC	David Black
	Erin Bournival
Hewlett Packard Enterprise	Jeff Wolford
	Erik Morales
	Curtis Ballard
	Wayne Bellamy
	Neil Wanamaker
HGST	Paul Suhler
	Chet Mercado
	David Brewer
	Frank Chu
	Joe Breher
Intel	James Boyd
	Steven Wells
Marvell	John Schadegg
	Kevin Tonthat
	Paul Wassenberg
Micron	John Goldman
	Daniel Hubbard
	Alan Haffner
	Carl Mies
	Jafar Naji
	Jim Cooke
	Michael George
	Neal Galbo
	Victor Tsai
NetApp	Fred Knight
	Tim Emami
OCZ Storage Solutions	Tom Friend
PMC-Sierra	Adnan Jiwani
	Sanjay Goyal

<b>Company</b>	<b>Name</b>
Samsung Semiconductor	William Martin
	KeunSoo Jo
SanDisk	Avraham Shimor
	David Landsman
	Yoni Shternhell
Seagate	Jim Hatfield
	Gerry Houlder
	Josh Bingaman
	Mary Dunn
	Tom Lenny
Toshiba	Patrick Hery
	Dan Colegrove
	Mark Carlson
	Mike Fitzpatick
Western Digital	Curtis Stevens
	Marvin DeForest
	Nathan Obr
	Ralph Weber

## Introduction

This standard encompasses the following:

Clause 1 describes the scope.

Clause 2 provides normative references for the entire standard.

Clause 3 provides definitions, abbreviations, and conventions used within the entire standard.

Clause 4 describes the feature sets defined by this standard.

Clause 5 describes commands.

Clause 6 describes logs.

Clause 7 describes command normal and error outputs.

Annex A describes zoned devices from a host's perspective.

American National Standard  
for Information Technology –

# Zoned Device ATA Command Set (ZAC)

## 1 Scope

The set of AT Attachment standards consists of this standard and the ATA implementation standards described in AT Attachment - 8 ATA/ATAPI Architecture Model (ATA8-AAM). This standard specifies the command set that host systems use to access storage devices that implement the Host Aware Zones feature set (see 4.3) or the Host Managed Zones feature set (see 4.4). This standard provides a common command set for systems manufacturers, system integrators, software suppliers, and suppliers of storage devices that provide one of the zones feature sets. Figure 1 shows the relationship of this standard to other ATA standards as well as related device and host standards and specifications (e.g., SCSI standards and SATA-IO specifications).

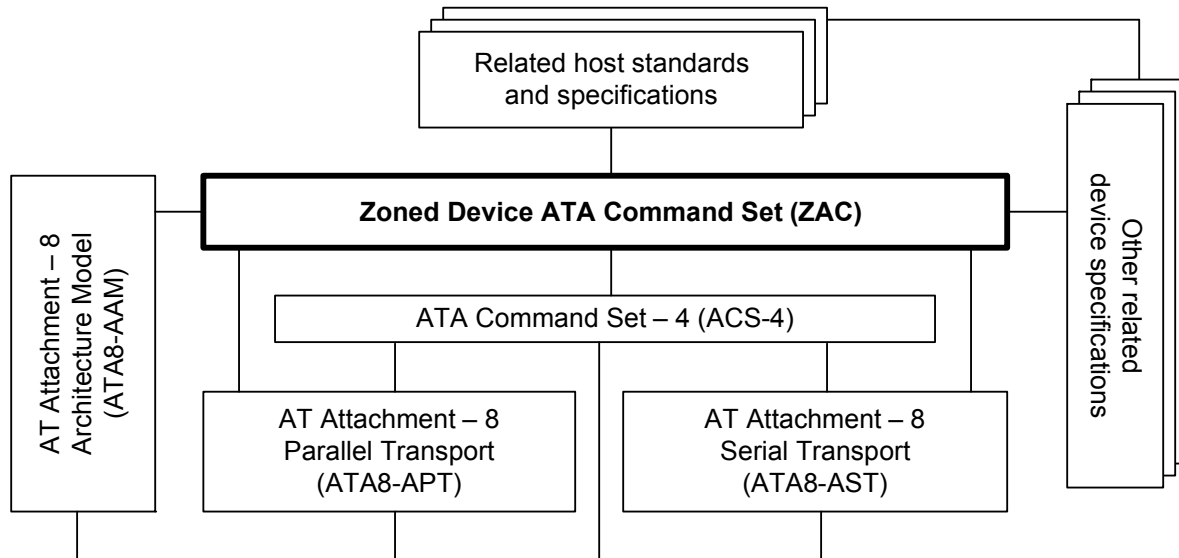


Figure 1 — ATA document relationships

This standard maintains compatibility with the ACS-4 standard, INCITS 529-201x, while providing additional functions.

## 2 Normative references

The following documents, in whole or in part, are normatively referenced in this document and are indispensable for its application. For dated references, only the edition cited applies. For undated references, the latest edition of the referenced document (including any amendments) applies.

ISO/IEC 646, *Information technology – ISO 7-bit coded character set for information interchange*<sup>1</sup>

ANSI INCITS 451-2008, *AT Attachment-8 – ATA/ATAPI Architecture Model (ATA8-AAM)*

ANSI INCITS 493-2012, *AT Attachment-8 – Serial Transport (ATA8-AST)*

T13/INCITS BSR 535, *ATA/ATAPI Serial Transport (AST-2)* (under consideration)

ANSI INCITS 524-201x, *AT Attachment-8 – Parallel Transport (ATA8-APT)* (under national consideration)

ANSI INCITS 522-2014, *ATA/ATAPI Command Set - 3 (ACS-3)* (planned as ISO/IEC 17760-103)

T13/INCITS BSR 529, *ATA Command Set - 4 (ACS-4)* (under consideration)

T10/INCITS BSR 502, *SCSI Primary Commands - 5 (SPC-5)* (under consideration)

Serial ATA revision 3.2 (SATA 3.2)<sup>2</sup>

Serial ATA revision 3.3 (SATA 3.3) (under consideration)<sup>2</sup>

---

<sup>1</sup> ANSI INCITS 4-1986 (R2012), *Information Systems – Coded Character Sets – 7-Bit American National Standard Code for Information Interchange (7-Bit ASCII)*

<sup>2</sup> For more information about Serial ATA specifications see <http://www.sata-io.org>.



## 3 Definitions, abbreviations, and conventions

### 3.1 Definitions

#### 3.1.1 28-bit command

command that uses the FEATURE field (7:0), COUNT field (7:0), LBA field (27:0), ICC field (7:0), AUXILIARY field (31:0), DEVICE field (15:8), and COMMAND field (7:0) to specify its arguments

Note 1 to entry: See ACS-4.

#### 3.1.2 48-bit command

command that uses the FEATURE field (15:0), COUNT field (15:0), LBA field (47:0), ICC field (7:0), AUXILIARY field (31:0), DEVICE field (15:8), and COMMAND field (7:0) to specify its arguments

Note 1 to entry: See ACS-4.

#### 3.1.3 accessible max address

maximum LBA that is accessible by read commands and write commands that return command completion without error

Note 1 to entry: See ACS-4.

#### 3.1.4 additional sense code

combination of the ADDITIONAL SENSE CODE field and the ADDITIONAL SENSE CODE QUALIFIER field

Note 1 to entry: See ACS-4.

#### 3.1.5 application client

object in the host that is the source of commands and device management functions

Note 1 to entry: See ATA8-AAM.

#### 3.1.6 ATA device

device that implements the General feature set

Note 1 to entry: See ACS-4.

#### 3.1.7 background activity

activity initiated by a command that occurs after command completion has been reported

#### 3.1.8 byte

sequence of eight contiguous bits considered as a unit

Note 1 to entry: See ACS-4.

#### 3.1.9 cache

data storage area outside the area accessible by hosts

Note 1 to entry: The cache may contain a subset of the data stored in the non-volatile media.

#### 3.1.10 command aborted

command completion with the ERROR bit set to one in the STATUS field and the ABORT bit set to one in the ERROR field

Note 1 to entry: See ACS-4.

#### 3.1.11 command acceptance

positive acknowledgement of a command being received by a device

Note 1 to entry: See ACS-4.

#### 3.1.12 command completion

completion by the device of processing in which the device:

- a) completes the action requested by the command or terminates the command with an error; and
- b) sets the appropriate bits in the ERROR field and the STATUS field

Note 1 to entry: See ACS-4.

#### 3.1.13 Conventional zone

zone that is not associated with a write pointer

Note 1 to entry: See 4.6.2.

**3.1.14 device**

data storage peripheral

Note 1 to entry: An example of a device is a disk drive.

**3.1.15 DWord**

sequence of four contiguous bytes considered as a unit

Note 1 to entry: See ACS-4.

**3.1.16 ending LBA**

highest numbered LBA of a read command or write command

**3.1.17 flush command**

command that flushes the volatile write cache

Note 1 to entry: See ACS-4.

**3.1.18 highest LBA**

highest numbered LBA of a range of LBAs

**3.1.19 host**

object that originates commands and device management functions

Note 1 to entry: See ATA8-AAM.

**3.1.20 LBA**

value used to reference a logical sector

**3.1.21 log**

named sequence of one or more log pages

Note 1 to entry: See ACS-4.

**3.1.22 log address**

numeric value that a log command (see ACS-4) uses to identify a specific log

Note 1 to entry: See ACS-4.

**3.1.23 log page**

512-byte block of data associated with a log

Note 1 to entry: See ACS-4.

**3.1.24 logical sector**

set of words accessed and referenced as a unit that contain user data and are referenced by LBA

Note 1 to entry: See ACS-4.

**3.1.25 lowest LBA**

lowest numbered LBA of a range of LBAs

**3.1.26 media**

material on which user data is stored

**3.1.27 native max address**

LBA that a device reports by GET NATIVE MAX ADDRESS EXT command (see ACS-4) and the maximum LBA accepted by a device using the SET ACCESSIBLE MAX ADDRESS EXT command (see ACS-4)

Note 1 to entry: See ACS-4.

**3.1.28 NCQ command**

command that is in the NCQ feature set

Note 1 to entry: See ACS-4.

**3.1.29 non-NCQ command**

command that is not in the NCQ feature set (see ACS-4)

**3.1.30 non-sequential write command**

write command processed on a write pointer zone for which the lowest LBA for that command is not equal to the write pointer for that write pointer zone

**3.1.31 non-volatile media**

physical storage media that retains user data written to it through all reset events (e.g., power-on reset)

**3.1.32 open zone**

zone with a Zone Condition of IMPLICITLY OPENED or EXPLICITLY OPENED

**3.1.33 open zone resources**

zone resources that are available only to an open zone

Note 1 to entry: The availability of open zone resources is determined by the processing of a Manage Open Zone Resources function (see 4.6.3.4.9).

Note 2 to entry: Open zone resources may be volatile.

**3.1.34 persistent zone resources**

zone resources that persist across a power-on reset

Note 1 to entry: For a zone in any Zone Condition other than IMPLICITLY OPENED or EXPLICITLY OPENED, all zone resources are persistent zone resources.

**3.1.35 physical sector**

one or more contiguous logical sectors that are read from or written to the device media in a single operation

Note 1 to entry: See ACS-4.

**3.1.36 power cycle**

when power is removed from a host or device until the subsequent power-on event

Note 1 to entry: See ACS-4.

**3.1.37 power-on reset**

host specific routine performed by the host or the routine performed by a device after detecting a power-on event

Note 1 to entry: See ACS-4.

**3.1.38 QWord**

sequence of eight contiguous bytes considered as a unit

Note 1 to entry: See ACS-4.

**3.1.39 read command**

command that causes the device to retrieve user data

Note 1 to entry: See ACS-4.

Note 2 to entry: This standard defines requirements (e.g., see 4.6.3.2.2 and 4.6.3.3.2) for read commands in addition to those defined in ACS-4.

**3.1.40 read FUA command**

read command that specifies that the device read user data from non-volatile media

Note 1 to entry: The READ FPDMA QUEUED command (see ACS-4) with the FUA bit set to one is a read FUA command.

**3.1.41 SATA device**

device implementing the serial ATA transport

Note 1 to entry: See ATA8-AST.

**3.1.42 sense data**

sense key and additional sense code

Note 1 to entry: See ACS-4.

**3.1.43 sense key**

contents of the SENSE KEY field in the sense data

Note 1 to entry: See ACS-4 and SPC-4.

**3.1.44 Sequential Write Preferred zone**

write pointer zone in which the device allows write commands that specify a starting LBA that is within the zone but is not equal to that zone's write pointer

Note 1 to entry: See 4.6.3.2.

**3.1.45 Sequential Write Required zone**

write pointer zone in which the device requires that write commands specify a starting LBA that is equal to the LBA indicated by that zone's write pointer

Note 1 to entry: See 4.6.3.3.

**3.1.46 signature**

unique set of values placed in the return parameters used to distinguish device types (e.g., ATA device)

Note 1 to entry: See ACS-4.

**3.1.47 starting LBA**

lowest numbered LBA of a read command or a write command

**3.1.48 state machine variable**

variable that exists within the context of a state machine

Note 1 to entry: A state machine variable may contain status from one state that is used in another state of the same state machine.

Note 2 to entry: The value contained in a state machine variable may affect subsequent state transitions or state machine outputs.

**3.1.49 stored data**

user data that has been stored by a write command

**3.1.50 substitute data pattern**

data that is substituted for data being read as user data if no user data is available

Note 1 to entry: The device shall store the data in the substitute data pattern by writing the data to non-volatile media.

Note 2 to entry: The substitute data pattern may be written as user data by the FINISH ZONE EXT command (see 5.2.5).

Note 3 to entry: The substitute data pattern for Sequential Write Preferred zones is described in 4.6.3.2.3.

Note 4 to entry: The substitute data pattern for Sequential Write Required zones is described in 4.6.3.3.3.

**3.1.51 transport**

mechanism used to communicate with a device

Note 1 to entry: See ATA8-APT and ATA8-AST.

**3.1.52 unrecoverable error**

error that causes the device to set either the ERROR bit or the DEVICE FAULT bit to one in the STATUS field at command completion

Note 1 to entry: See ACS-4.

**3.1.53 user data**

data that is transferred between the application client and the device using read commands and write commands

**3.1.54 volatile cache**

cache that does not retain data through power cycles

**3.1.55 word**

sequence of two contiguous bytes considered as a unit

Note 1 to entry: See ACS-4.

**3.1.56 write command**

command that causes the device to store user data

Note 1 to entry: See ACS-4.

Note 2 to entry: This standard defines requirements (e.g., see 4.6.3.2.1 and 4.6.3.3.1) for write commands in addition to those defined in ACS-4.

**3.1.57 write FUA command**

write command that specifies that the device transfer user data from the host to non-volatile media before reporting command completion

Note 1 to entry: The WRITE FPDMA QUEUED command (see ACS-4) with the FUA bit set to one, the WRITE MULTIPLE FUA EXT command (see ACS-4), and the WRITE DMA FUA EXT command (see ACS-4) are write FUA commands.

**3.1.58 write pointer**

pointer to a logical sector in a write pointer zone where the next write command in that zone should start

Note 1 to entry: See 4.6.3.1.

**3.1.59 write pointer zone**

zone that has an associated write pointer

Note 1 to entry: See 4.6.3.

**3.1.60 zone**

one of a set of disjoint contiguous ranges of LBAs that comprise a device

Note 1 to entry: See 4.2.

**3.1.61 zone attribute**

named property associated with a zone

Note 1 to entry: In the absence of other conditions or actions (e.g., actions by the host), zone attributes do not affect the operation of the Zone Condition state machine.

Note 2 to entry: See 4.5.

**3.1.62 zone resources**

device resources that are used by a zone

Note 1 to entry: This standard defines open zone resources and persistent zone resources.

Note 2 to entry: Zone resources include zone attributes.

**3.1.63 zoned device**

device that supports the Host Aware Zones feature set (see 4.3) or the Host Managed Zones feature set (see 4.4)

**3.2 Symbols and abbreviations****3.2.1 Abbreviations**

Abbreviations used in this standard:

<b>Abbreviation</b>	<b>Meaning</b>
ACS	ATA Command Set
ASCII	American Standard Code for Information Interchange
ATA	AT Attachment
ATA8-AAM	AT Attachment-8 - ATA/ATAPI Architecture Model
ATA8-APT	AT Attachment-8 - Parallel Transport
ATA8-AST	AT Attachment-8 - Serial Transport
DMA	Direct Memory Access
EXT	Command that uses the extended (48-bit LBA) format parameters
ISO	Organization for International Standards
LBA	Logical Block Address
NCQ	Native Command Queueing

<b>Abbreviation</b>	<b>Meaning</b>
SATA	Serial ATA
SATA-IO	Serial ATA International Organization ( <a href="http://www.sata-io.org">http://www.sata-io.org</a> )
SSP	Software Settings Preservation
ZAC	Zoned Device ATA Command Set

### 3.2.2 Units

Units used in this standard:

<b>Unit</b>	<b>Meaning</b>
min	minute (unit of time)
ms	millisecond (i.e., $10^{-3}$ seconds)
ns	nanosecond (i.e., $10^{-9}$ seconds)
s	second (unit of time)
$\mu$ s	microsecond (i.e., $10^{-6}$ seconds)

### 3.2.3 Symbols

Symbols used in this standard:

<b>Symbol</b>	<b>Meaning</b>
©	copyright
®	registered trademark

### 3.2.4 Mathematical operators

Mathematical operators used in this standard:

<b>Mathematical Operator</b>	<b>Meaning</b>
+	added to
×	multiplied by
/	divided by
<	less than
≤	less than or equal to
>	greater than
≥	greater than or equal to
AND	logical AND

## 3.3 Keywords

### 3.3.1 expected

keyword used to describe the behavior of the hardware or software in the design models assumed by this standard

Note 1 to entry: Other hardware and software design models may also be implemented.

### 3.3.2 mandatory

keyword indicating items to be implemented as defined by this standard

### 3.3.3 may

keyword that indicates flexibility of choice with no implied preference

### 3.3.4 N/A

keyword that indicates a field is not applicable and has no defined value

Note 1 to entry: A field defined as N/A should not be checked by the host or device.

**3.3.5 obsolete**

keyword indicating that the designated bits, bytes, words, fields, and code values that may have been defined in previous standards are not defined in this standard and shall not be reclaimed for other uses in future standards

Note 1 to entry: Some degree of functionality may be required for items designated as “obsolete” to provide for backward compatibility.

Note 2 to entry: Obsolete commands should not be used by the host.

Note 3 to entry: Commands defined as obsolete may return command aborted by devices conforming to this standard. However, if a device does not return command aborted for an obsolete command, then the device shall return command completion for the command.

**3.3.6 optional**

keyword that describes features that are not required by this standard

Note 1 to entry: If any optional feature defined by the standard is implemented, the feature shall be implemented in the way defined by the standard.

**3.3.7 prohibited**

keyword indicating that an item shall not be supported by an implementation

**3.3.8 reserved**

keyword indicating reserved bits, bytes, words, fields, and code values that are set aside for future standardization

Note 1 to entry: The use and interpretation of reserved bits, bytes, words, fields, and code values may be specified by future extensions to this or other standards.

Note 2 to entry: A reserved bit, byte, word, or field shall be cleared to zero, or set in accordance with a future extension to this standard.

Note 3 to entry: The recipient shall not check reserved bits, bytes, words, or fields.

Note 4 to entry: Receipt of reserved code values in defined fields shall be considered a command parameter error and reported by returning command aborted.

**3.3.9 retired**

keyword indicating that the designated bits, bytes, words, fields, and code values that had been defined in previous standards are not defined in this standard and may be reclaimed for other uses in future standards

Note 1 to entry: If retired bits, bytes, words, fields, or code values are used before they are reclaimed, they shall have the meaning or functionality as described in previous standards.

**3.3.10 shall**

keyword indicating a mandatory requirement

Note 1 to entry: Designers are required to implement all such mandatory requirements to ensure interoperability with other products that conform to this standard.

**3.3.11 should**

keyword indicating flexibility of choice with a strongly preferred alternative

Note 1 to entry: This keyword is equivalent to the phrase “it is recommended”.

**3.3.12 vendor specific**

something (e.g., a bit, field, code value) that is not defined by this standard

Note 1 to entry: Specification of the referenced item is determined by the device vendor and may be used differently in various implementations.

**3.4 Conventions****3.4.1 Overview**

Lowercase is used for words having the normal English language meaning. Certain words and terms used in this standard have a specific meaning beyond the normal English language meaning. These words and terms are defined either in clause 3 or in the text where they first appear.

Quantities having a defined numeric value are identified by large uppercase letters (e.g., the command code in the IDENTIFY DEVICE command).

Quantities having a discrete but unspecified value are identified using small uppercase letters (e.g., SEQUENTIAL WRITE PREFERRED is one possible value of the Zone Type zone attribute). Such quantities are associated with an observable behavior or value that is specific to a given implementation standard.

Names of fields are in small uppercase (e.g., COUNT). Normal case is used when the contents of a field are being discussed. Fields containing only one bit are usually referred to as the NAME bit instead of the NAME field. (See 3.4.5 for the naming convention used for numbering bits.)

The expression “word n” or “bit n” indicates the content of word n or the content of bit n.

### 3.4.2 Precedence

If there is a conflict between text, figures, and tables, the precedence shall be tables, figures, then text.

### 3.4.3 Lists

#### 3.4.3.1 Lists overview

Lists are associated with an introductory paragraph or phrase, and are numbered relative to that paragraph or phrase (i.e., all lists begin with an a) or 1) entry).

Each item in a list is preceded by an identification with the style of the identification being determined by whether the list is intended to be an ordered list or an unordered list.

If the item in a list is not a complete sentence, the first word in the item is not capitalized. If the item in a list is a complete sentence, the first word in the item is capitalized.

Each item in a list ends with a semicolon, except the last item, which ends in a period. The next to the last entry in the list ends with a semicolon followed by an “and” or an “or” (i.e., “...; and”, or “...; or”). The “and” is used if all the items in the list are required. The “or” is used if only one or more items in the list are required.

#### 3.4.3.2 Unordered lists

An unordered list is one in which the order of the listed items is unimportant (i.e., it does not matter where in the list an item occurs as all items have equal importance). Each list item shall start with a lower case letter followed by a close parenthesis. If it is necessary to subdivide a list item further with an additional unordered list (i.e., have a nested unordered list), then the nested unordered list shall be indented and each item in the nested unordered list shall start with an upper case letter followed by a close parenthesis.

The following is an example of an unordered list with a nested unordered list:

EXAMPLE - The following are the items for the assembly:

- a) a box containing:
  - A) a bolt;
  - B) a nut; and
  - C) a washer;
- b) a screwdriver; and
- c) a wrench.

#### 3.4.3.3 Ordered lists

An ordered list is one in which the order of the listed items is important (i.e., item n is required before item n+1). Each listed item starts with a Western-Arabic numeral followed by a close parenthesis. If it is necessary to subdivide a list item further with an additional unordered list (i.e., have a nested unordered list), then the nested unordered list shall be indented and each item in the nested unordered list shall start with an upper case letter followed by a close parenthesis.

The following is an example of an ordered list with a nested unordered list:

EXAMPLE - The following are the instructions for the assembly:

- 1) remove the contents from the box;
- 2) assemble the item;



- A) use a screwdriver to tighten the screws; and
- B) use a wrench to tighten the bolts;
- and
- 3) take a break.

### 3.4.4 Numbering

A binary number is represented in this standard by any sequence of digits consisting of only the Western-Arabic numerals 0 and 1 immediately followed by a lower-case b (e.g., 0101b). Underscores or spaces may be included between characters in binary number representations to increase readability or delineate field boundaries (e.g., 0 0101 1010b or 0\_0101\_1010b).

A hexadecimal number is represented in this standard by any sequence of digits consisting of only the Western-Arabic numerals 0 through 9 and/or the upper-case English letters A through F immediately followed by a lower-case h (e.g., FA23h). Underscores or spaces may be included between characters in hexadecimal number representations to increase readability or delineate field boundaries (e.g., B FD8C FA23h or B\_FD8C\_FA23h).

A decimal number is represented in this standard by any sequence of digits consisting of only the Arabic numerals 0 through 9 not immediately followed by a lower-case b or lower-case h (e.g., 25). This standard uses the following conventions for representing decimal numbers:

- a) the decimal separator (i.e., separating the integer and fractional portions of the number) is a period;
- b) the thousands separator (i.e., separating groups of three digits in a portion of the number) is a space;
- and
- c) the thousands separator is used in both the integer portion and the fraction portion of a number.

Table 1 shows some examples of decimal numbers using various numbering conventions.

**Table 1 — Numbering conventions**

French	English	This standard
0,6	0.6	0.6
3,141 592 65	3.14159265	3.141 592 65
1 000	1,000	1 000
1 323 462,95	1,323,462.95	1 323 462.95

A decimal number represented in this standard with an overline over one or more digits following the decimal point is a number where the overlined digits are infinitely repeating (e.g.,  $666.\overline{6}$  means  $666.666\ 666\dots$  or  $666\ 2/3$ , and  $12.\overline{142\ 857}$  means  $12.142\ 857\ 142\ 857\dots$  or  $12\ 1/7$ ).

### 3.4.5 Bit conventions


Name (n:m), where n is greater than m, denotes a set of bits (e.g., Feature (7:0)). n:m where n is greater than m denotes a bit range in a table.

### 3.4.6 Number range convention

The verbal form “p..q”, where p is less than q, represents a range of numbers (e.g., words 100..103 represents words 100, 101, 102, and 103).

**3.4.7 State diagram convention**

All state diagrams in this standard use the notation shown in figure 2.



**Figure 2 — State diagram convention**

Each state is identified by a state designator and a state name. The state designator is unique among all states in all state diagrams in this standard. The state designator consists of a set of letters that are capitalized in the title of the figure containing the state diagram followed by a unique number. The state name is a brief description of the purpose of the state, and the same state name may appear in more than one state diagram. The reference below that state name is a link to the subclause that contains a full definition of the state.

Transitions between states are indicated by lines with arrows. Each transition includes a transition reference to the paragraph or paragraphs that contains a full definition of the transition. If the transition exits a figure or enters a figure, the transition line goes to or from a state designator with double underlines rather than to or from a state.

**3.4.8 Offset convention**

An offset is a byte value used as an index into a larger data structure.

## 4 Feature set definitions

### 4.1 Introduction

Based on the basic zoned device model (see 4.2), the following feature sets are defined by this standard:

- a) Host Aware Zones feature set (see 4.3); and
- b) Host Managed Zones feature set (see 4.4).

### 4.2 Basic zoned device model

Zoned devices are accessed using LBAs. The LBAs are divided into ranges called zones. The entire capacity of a zoned device is organized into a set of contiguous, non-overlapping zones. Figure 3 shows a zoned device with  $n$  zones and  $m$  LBAs where LBA 0 is the lowest LBA of zone 0 and LBA  $m-1$  is the highest LBA of zone  $n-1$ .

<b>Zone</b>	0		1	...	$n-1$	
<b>LBA</b>	0	1	...			$m-1$

**Figure 3 — Zones in a zoned device**

Each zone is one of the following types:

- a) a Conventional zone (see 4.6.2); or
- b) a write pointer zone (see 4.6.3) that is:
  - A) a Sequential Write Preferred zone (see 4.6.3.2); or
  - B) a Sequential Write Required zone (see 4.6.3.3).

The REPORT ZONES EXT command (see 5.2.7) returns the zone type for each zone.

### 4.3 Host Aware Zones feature set

Devices that support the Host Aware Zones feature set shall support all feature sets that are defined as mandatory by ACS-4, shall not support feature sets that are defined as prohibited by ACS-4, and may support feature sets that are defined as optional by ACS-4.

Devices that support the Host Aware Zones feature set:

- a) shall report the ATA device signature (see ACS-4);
- b) shall support at least one Sequential Write Preferred zone (see 4.6.3.2);
- c) may support one or more Conventional zones (see 4.6.2);
- d) shall not support any Sequential Write Required zones (see 4.6.3.3);
- e) shall support the Zoned Device Information page (see 6.2.2) in the IDENTIFY DEVICE data log (see ACS-4);
- f) shall set the ZONED field to 01b in the Supported Capabilities page of the IDENTIFY DEVICE data log (see ACS-4);
- g) shall support the Sense Data Reporting feature set (see ACS-4);
- h) shall support the SCT Write Same command (see ACS-4 and 4.7.9);
- i) shall support the non-NCQ command format for those commands defined as mandatory in table 13; and
- j) if the NCQ feature set (see ACS-4) is supported, then:
  - A) shall support the NCQ command encapsulations (see 5.2.2, 5.2.3, and ACS-4) for those commands defined as mandatory in table 13;
  - B) shall support NCQ autosense (i.e., set the NCQ AUTOSENSE SUPPORTED bit to one in the Serial ATA page of the IDENTIFY DEVICE data log (see ACS-4));
  - C) shall support bytes 512..1 023 of the NCQ Command Error log (see ACS-4); and
  - D) shall support the use of the NCQ command encapsulations of GPL feature set commands (see ACS-4) to read and write:
    - a) the SCT Command/Status log (see ACS-4); and

- b) the SCT Data Transfer log (see ACS-4).

Devices that support the Host Aware Zones feature set shall support the applicable features described in 4.7.

#### 4.4 Host Managed Zones feature set

The Host Managed Zones feature set relies on active host management of interactions with the device to maintain efficient operation of the device.

Devices that support the Host Managed Zones feature set shall support the following feature sets defined in ACS-4:

- a) the 48-bit Address feature set;
- b) the General feature set;
- c) the General Purpose Logging (GPL) feature set;
- d) the Native Command Queuing (NCQ) feature set, if the device is a SATA device; and
- e) the Power Management feature set.

Devices that support the Host Managed Zones feature set:

- a) shall report the Host Managed Zoned device signature (see ACS-4);
- b) shall support at least one Sequential Write Required zone (see 4.6.3.3);
- c) may support one or more Conventional zones (see 4.6.2);
- d) shall not support any Sequential Write Preferred zones (see 4.6.3.2);
- e) shall support the Zoned Device Information page (see 6.2.2) in the IDENTIFY DEVICE data log (see ACS-4);
- f) shall set the ZONED field to 00b in the Supported Capabilities page of the IDENTIFY DEVICE data log (see ACS-4);
- g) shall support the Sense Data Reporting feature set (see ACS-4) and shall:
  - A) enable the Sense Data Reporting feature set during the processing of any reset; and
  - B) process the SET FEATURES Enable/Disable the Sense Data Reporting feature set subcommand as defined in ACS-4;
- h) shall support the SCT Write Same command (see ACS-4 and 4.7.9);
- i) shall support the non-NCQ command format for those commands defined as mandatory in table 13; and
- j) if the NCQ feature set (see ACS-4) is supported, then:
  - A) shall support the NCQ command encapsulations (see 5.2.2, 5.2.3, and ACS-4) for those commands defined as mandatory in table 13;
  - B) shall support NCQ autosense (i.e., set the NCQ AUTOSENSE SUPPORTED bit to one in the Serial ATA page of the IDENTIFY DEVICE data log (see ACS-4));
  - C) shall support bytes 512..1 023 of the NCQ Command Error log (see ACS-4); and
  - D) shall support the use of the NCQ command encapsulations of GPL feature set commands (see ACS-4) to read and write:
    - a) the SCT Command/Status log (see ACS-4); and
    - b) the SCT Data Transfer log (see ACS-4).

If the device supports the Self-Monitoring, Analysis, and Reporting Technology (SMART) feature set, then the device:

- a) shall clear the SELECTIVE SELF-TEST IMPLEMENTED bit to zero in the Off-line data collection capabilities byte of the SMART data structure (see ACS-3); and
- b) shall not support the Selective Self-Test log (see ACS-3).

Devices that support the Host Managed Zones feature set shall not support the following feature sets:

- a) the PACKET feature set (see ACS-3);
- b) the Accessible Max Address Configuration feature set (see ACS-4);
- c) the CompactFlash Association (CFA) feature set (see ACS-4);
- d) the Security feature set; and
- e) the Streaming feature set (see ACS-4).

Devices that support the Host Managed Zones feature set shall support the applicable features described in 4.7.

## 4.5 Zone attributes

### 4.5.1 Summary

The zoned device zone attributes are summarized in table 2.

**Table 2 — Summary of zone attributes**

Attribute	Description	Reference
Zone Type	Type of zone (e.g., conventional)	4.5.2
Zone Condition	Zone's operational characteristics	4.5.3
WPointer	Combination of the write pointer (see 4.6.3.1) and whether the write pointer is valid	4.5.4
RWP Recommended	Indication of whether the processing of a RESET WRITE POINTER EXT command (see 5.2.8) is recommended	4.5.5
Non-Sequential Write Resources Active	Indication of whether zone resources are allocated for non-sequential writes	4.5.6

### 4.5.2 Zone Type zone attribute

Zone Type indicates (see table 3) the type of zone.

**Table 3 — Zone Type zone attribute**

Zone Type	Reference
CONVENTIONAL	4.6.2
SEQUENTIAL WRITE PREFERRED	4.6.3.2
SEQUENTIAL WRITE REQUIRED	4.6.3.3

The relationships between Zone Type and other zone attributes are shown in table 4.

**Table 4 — Relationships between zone attributes**

Zone Type	Other zone attributes			
	Non-Sequential Write Resources Active	RWP Recommended	WPointer	Zone Condition
CONVENTIONAL	false <sup>a</sup>	false <sup>a</sup>	invalid <sup>a</sup>	see 4.5.3
SEQUENTIAL WRITE PREFERRED	see 4.5.6	see 4.5.5	see 4.5.4	
SEQUENTIAL WRITE REQUIRED	false <sup>b</sup>		and table 7	
<sup>a</sup> The zone attributes values for a zone with Zone Type CONVENTIONAL are required to be compatible with devices that conform to ACS-4. <sup>b</sup> Zone resources for non-sequential writes are not allocated for a zone in which all non-sequential write commands return command completion with an error.				

#### 4.5.3 Zone Condition zone attribute

The Zone Condition (see table 5) is determined by the Zone Type (see 4.5.2) and the Zone Condition state machine state (see 4.6.3.4).

**Table 5 — Zone Condition zone attribute**

Zone Type	Zone Condition state machine state	Reference	Zone Condition
CONVENTIONAL	outside the scope of this standard	4.6.2	NOT WRITE POINTER
SEQUENTIAL WRITE PREFERRED or SEQUENTIAL WRITE REQUIRED	ZC1: Empty state	4.6.3.4.2	EMPTY
	ZC2: Implicit_Open state	4.6.3.4.3	IMPLICITLY OPENED
	ZC3: Explicit_Open state	4.6.3.4.4	EXPLICITLY OPENED
	ZC4: Closed state	4.6.3.4.5	CLOSED
	ZC5: Full state	4.6.3.4.6	FULL
	ZC6: Read_Only state	4.6.3.4.7	READ ONLY
	ZC7: Offline state	4.6.3.4.8	OFFLINE

#### 4.5.4 WPointer zone attribute

For write pointer zones (see 4.6.3), WPointer is a combination of the write pointer (see 4.6.3.1) and whether the write pointer is valid (see table 7).

If Zone Type is CONVENTIONAL, WPointer shall be invalid.

#### 4.5.5 RWP Recommended zone attribute

For write pointer zones (see 4.6.3), if RWP Recommended is:

- a) true, then the device has determined (e.g., through the detection of insufficient zone resources) that the processing of a RESET WRITE POINTER EXT command (see 5.2.8) specifying this zone is recommended; or
- b) false, then the device has no recommendation for or against the processing of a RESET WRITE POINTER EXT command specifying this zone.

RWP Recommended shall be set to false if:

- a) the Zone Type is CONVENTIONAL; or
- b) the Zone Condition is EMPTY, READ ONLY, or OFFLINE.

The device shall set the sense key to COMPLETED and the additional sense code set to ZONE RESET WRITE POINTER RECOMMENDED (see ACS-4 and SPC-5), if:

- a) a command completes without an error; and
- b) the RWP Recommended attribute in at least one zone has changed from false to true since the last time a command completed without an error.

#### 4.5.6 Non-Sequential Write Resources Active zone attribute

For Sequential Write Preferred zones (see 4.6.3.2), if Non-Sequential Write Resources Active is:

- a) true, then the device has zone resources allocated for the processing of non-sequential writes in this zone; or
- b) false, then the device does not have zone resources allocated for the processing of non-sequential writes in this zone.

If a non-sequential write command is processed in a zone, the device may set Non-Sequential Write Resources Active to true. The device may set Non-Sequential Resources Active zone to false at any time if:

- a) the write pointer does not indicate the lowest LBA in the zone; and
- b) zone resources are no longer allocated for the processing of non-sequential writes in this zone.

If the write pointer indicates the lowest LBA in the zone and the device has processed any non-sequential write commands in a zone since the last time the Zone Condition was EMPTY, then Non-Sequential Resources Active shall not be set to false until the write pointer does not indicate the lowest LBA in the zone.

The Non-Sequential Write Resources Active shall be set to false if:

- a) the Zone Type is CONVENTIONAL or SEQUENTIAL WRITE REQUIRED;
- b) the Zone Condition is EMPTY, READ ONLY, or OFFLINE; or
- c) the device has not processed a non-sequential write command in a zone since the last time the Zone Condition was EMPTY.

## 4.6 Zone types

### 4.6.1 Overview

In a zoned block device, each zone:

- a) has the zone attributes defined in 4.5; and
- b) is one of the following zone types:
  - A) a Conventional zone (see 4.6.2); or
  - B) a write pointer zone (see 4.6.3) that is either:
    - a) a Sequential Write Preferred zone (see 4.6.3.2); or
    - b) a Sequential Write Required zone (see 4.6.3.3).

### 4.6.2 Conventional zones

A Conventional zone is a type of zone:

- a) that is not associated with a write pointer; and
- b) for which the device processes commands as described in ACS-4.

### 4.6.3 Write pointer zones

#### 4.6.3.1 Features common to all write pointer zones

##### 4.6.3.1.1 General


A write pointer zone is a type of zone that maintains a write pointer that indicates a logical sector within that zone. The write pointer may be updated as a result of processing of a write command. Depending on the state of a zone, the write pointer for that zone may or may not be valid (see 4.6.3.4.1).

Each write pointer zone has one write pointer. That write pointer indicates the LBA that the host should specify in the starting LBA of the next write command to that zone. Some commands (e.g., the RESET WRITE POINTER EXT command (see 5.2.8)) set the write pointer to the lowest LBA in the specified zone.

Each write pointer zone in a device is in one of the states defined by the Zone Condition state machine (see 4.6.3.4). The state of each zone is indicated by the Zone Condition (see 4.5.3) and affects the contents of the ZONE CONDITION field in the data returned by the REPORT ZONES EXT command (see 5.2.7).

Read commands that complete without an error do not affect the write pointer.

Figure 4 shows a write pointer zone where no LBAs have been written since the most recent time the zone's Zone Condition was EMPTY with the write pointer set to the lowest LBA of that zone.

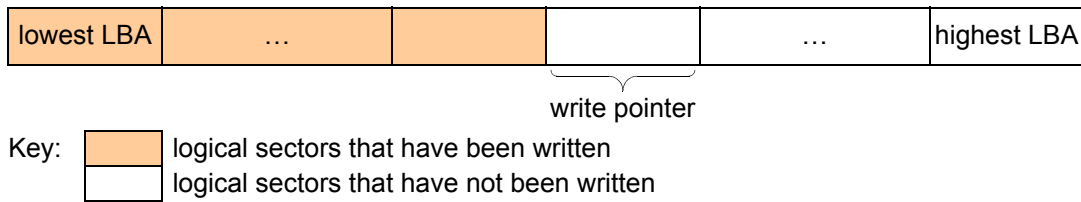


Key:  logical sectors that have not been written

**Figure 4 — Write pointer zone and write pointer after Reset Write Pointer function with no subsequent writes**

Figure 5 shows a write pointer zone where:

- a) some LBAs have been written since the most recent time the zone's Zone Condition was EMPTY (i.e., LBAs less than the write pointer value); and
- b) other LBAs have not been written since the most recent time the zone's Zone Condition was EMPTY (i.e., LBAs greater than or equal to the write pointer value).




**Figure 5 — Write pointer zone and write pointer example**

The example in figure 6 shows the effects of host actions on the write pointer and user data as follows:


- 1) at time 1, the write pointer zone starts in Zone Condition EMPTY with the write pointer set to the lowest LBA in that zone;
- 2) at time 2, a write command:
  - 1) changes the zone to Zone Condition IMPLICITLY OPENED; and
  - 2) writes user data starting at that zone's lowest LBA and sets the write pointer to the ending LBA written by the command plus one;
- 3) at time 3, a write command:
  - 1) writes user data that starts at the write pointer after time 2 and extends to the end of that zone;
  - 2) changes the zone to Zone Condition FULL; and
  - 3) changes the write pointer to invalid;
 and
- 4) at time 4, a RESET WRITE POINTER EXT command (see 5.2.8):
  - 1) sets the write pointer to the lowest LBA of that write pointer zone; and
  - 2) changes the zone to Zone Condition EMPTY.




**Time 1** – The write pointer is set to the lowest LBA of that zone.




**Time 2** – A write command increases the write pointer.



**Time 3** – A write command writes user data to the end of that zone and the write pointer becomes invalid.



**Time 4** – A RESET WRITE POINTER EXT command sets the write pointer back to the lowest LBA of that zone.



Key:  logical sectors that have been written  
 logical sectors that have not been written

**Figure 6 — Example effects of host actions on a write pointer zone**

**4.6.3.1.2 Resetting the write pointer**

A RESET WRITE POINTER EXT command (see 5.2.8) sets the write pointer in a write pointer zone to the lowest LBA of that zone and changes the Zone Condition to EMPTY.

The following ACS-4 feature sets and commands may set the write pointer in a write pointer zone to the lowest LBA of that zone and change the Zone Condition to EMPTY:

- a) the Sanitize Device feature set (see 4.7.6);
- b) the SECURITY ERASE UNIT command (see 4.7.8); and
- c) the SCT Write Same command (see 4.7.9).

**4.6.3.2 Sequential Write Preferred zones**

**4.6.3.2.1 Writing in Sequential Write Preferred zones**

A Sequential Write Preferred zone is a write pointer zone (see 4.6.3.1) in which the device:

- a) allows write commands that specify a starting LBA that is not equal to that zone's write pointer (i.e., non-sequential writes); and
- b) may allocate zone resources for the processing of non-sequential writes as described in 4.5.6.

A write command that starts at the write pointer of a Sequential Write Preferred zone and ends at the highest LBA of that Sequential Write Preferred zone results in the Zone Condition changing to FULL.

Figure 7 shows an example of a write command that starts at the write pointer associated with a Sequential Write Preferred zone.




Key:  logical sectors written by the write command  
 logical sectors not written by the write command


**Figure 7 — Example write command that starts at the write pointer**

Figure 8 shows examples of write commands that do not start at the write pointer (i.e., non-sequential write commands).


**Example write command that starts lower and ends at the initial write pointer**



**Example write command that starts lower and ends higher than the initial write pointer**



**Example write command that starts higher than the initial write pointer**



Key:  logical sectors written by the write command  
 logical sectors not written by the write command

**Figure 8 — Examples of write commands that do not start at the write pointer**

A non-sequential write command shall result in:

- a) the associated zone's write pointer being set to an LBA that is:
  - A) greater than or equal to the value of that write pointer before that command was processed; and
  - B) less than or equal to the minimum of:
    - a) one plus the highest LBA written in the zone since the last time the Zone Condition was EMPTY; and
    - b) the highest LBA in the zone;
- or
- b) the Zone Condition becoming FULL.

The device may incur delays during the processing of subsequent commands if completing a command results in:

- a) the number of Sequential Write Preferred zones that are open zones exceeding the value in the OPTIMAL NUMBER OF OPEN SEQUENTIAL WRITE PREFERRED ZONES field (see 6.2.2.3); or

- b) the number of Sequential Write Preferred zones in which Non-Sequential Write Resources Active is true to exceed the value in the OPTIMAL NUMBER OF NON-SEQUENTIALLY WRITTEN SEQUENTIAL WRITE PREFERRED ZONES field (see 6.2.2.4).

As described in the Zone Condition state machine (see 4.6.3.4), a non-sequential write command in a Sequential Write Preferred zone may result in the Zone Condition becoming FULL. If a non-sequential write command completes without error and results in the Zone Condition becoming FULL, then the device shall set the sense key to COMPLETED and the additional sense code set to ZONE TRANSITION TO FULL (see ACS-4 and SPC-4).

Regardless of the starting LBA in a write command, the resulting value of the write pointer is outside the scope of this standard if:

- a) at least one logical sector has been written;
- b) not all the logical sectors specified by the command have been written; and
- c) volatile write caching:
  - A) is not enabled and the write command returns command completion with an error; or
  - B) is enabled and an error occurs after the command returns command completion without an error.

#### 4.6.3.2.2 Reading in Sequential Write Preferred zones

A read command that specifies a logical sector that has been written since the most recent time the Zone Condition was EMPTY in a Sequential Write Preferred zone shall use the most recent stored data when processing the command. If the device processes a read command that specifies logical sectors that have not been written since the most recent time the Zone Condition was EMPTY, then for each logical sector that has not been written since the most recent time the Zone Condition was EMPTY the device shall:

- a) use the substitute data pattern for Sequential Write Preferred zones (see 4.6.3.2.3) as user data; and
- b) not write the substitute data pattern if a read FUA command is being processed.

#### 4.6.3.2.3 Substitute data pattern for Sequential Write Preferred zones

For Sequential Write Preferred zones, the substitute data pattern is data that shall be set by:

- a) the manufacturer; or
- b) the most recent successful:
  - A) OVERWRITE EXT command (see ACS-4), if any;
  - B) SECURITY ERASE UNIT command with the ERASE MODE bit cleared to zero (see ACS-4), if any;
 or
  - C) SCT Write Same command (see ACS-4), if any, with:
    - a) the FUNCTION CODE field set to 0001h, 0002h, 0101h, or 0102h; and
    - b) the WRITE AND SET SUBSTITUTE DATA PATTERN bit (see 4.7.9.2) set to one.

#### 4.6.3.3 Sequential Write Required zones

##### 4.6.3.3.1 Writing in Sequential Write Required zones

A Sequential Write Required zone is a write pointer zone (see 4.6.3.1) in which the device requires that write commands specify a starting LBA that is equal to that zone's write pointer.

If the device processes a write command without error, the write pointer is modified to indicate the LBA where a subsequent write command within the Sequential Write Required zone is able to be performed.

If the ending LBA of a write command that completed without error is equal to the highest LBA of a Sequential Write Required zone, then the Zone Condition becomes FULL and the write pointer becomes invalid.

If a write command returns command completion with an error, the value of the write pointer may be unknown to the host. The device indicates the current write pointer in at least one of the following places:

- a) the NCQ Command Error log (see ACS-4);
- b) the LBA field of the error outputs (see 4.6.3.3.1 and 4.6.3.3.2); and
- c) the data returned by the REPORT ZONES EXT command (see 5.2.7).

If the device processes a write command with:

- a) the starting LBA equal to the write pointer;

- b) the ending LBA is in the same Sequential Write Required zone; and
- c) an ending LBA that is equal to the last logical sector within a physical sector (see ACS-4),

then the device shall process that write command.

If the device processes a write command with:

- a) the starting LBA in a Sequential Write Required zone that is not equal to the write pointer; or
- b) an ending LBA that is not equal to the last logical sector within a physical sector (see ACS-4),

then the device shall:

- a) not write any data;
- b) return command completion with an error; and
- c) set the sense key to ILLEGAL REQUEST and the additional sense code to UNALIGNED WRITE COMMAND (see ACS-4 and SPC-5).

If the device processes a write command with an ending LBA that is not in the same Sequential Write Required zone as the starting LBA, then the device shall:

- a) not write any data;
- b) return command completion with an error; and
- c) set the sense key to ILLEGAL REQUEST and the additional sense code to WRITE BOUNDARY VIOLATION (see ACS-4 and SPC-5).

If the device processes a write command that specifies:

- a) a starting LBA that is not in a Sequential Write Required zone; and
- b) access to an LBA that is in a Sequential Write Required zone,

then the device shall:

- a) not write any data;
- b) return command completion with an error; and
- c) set the sense key to ILLEGAL REQUEST and the additional sense code to WRITE BOUNDARY VIOLATION (see ACS-4 and SPC-5).

If the device sets the additional sense code to UNALIGNED WRITE COMMAND or WRITE BOUNDARY VIOLATION, then the device shall set the LBA field in the error outputs to the value of the write pointer associated with the Sequential Write Required zone specified by the starting LBA in the command.

If a write command is not able to write all of the specified logical sectors before encountering an error that prevents further writing to the media, then the resulting value of the write pointer is outside the scope of this standard. The host should use the REPORT ZONES EXT command (see 5.2.7) to determine the current write pointer and other characteristics of the zone.

#### 4.6.3.3.2 Reading in Sequential Write Required zones

If the device processes a read command with:

- a) the starting LBA in a Sequential Write Required zone; and
- b) the ending LBA:
  - A) in the same Sequential Write Required zone; and
  - B) less than the write pointer,

then the device shall process that read command.

If the URSWRZ bit (see 6.2.2.2) is cleared to zero and device processes a read command with:

- a) the starting LBA:
  - A) less than the write pointer and the ending LBA greater than or equal to the write pointer in the same zone; or
  - B) greater than or equal to the write pointer and the ending LBA in the same zone,
 then the device shall:
  - A) not read any data;
  - B) return command completion with an error; and

- C) set the sense key to ILLEGAL REQUEST and the additional sense code to ATTEMPT TO READ INVALID DATA (see ACS-4 and SPC-5);
- or
- b) the ending LBA in a different Sequential Write Required zone than the starting LBA, then the device shall:
  - A) not read any data;
  - B) return command completion with an error; and
  - C) set the sense key to ILLEGAL REQUEST and the additional sense code to READ BOUNDARY VIOLATION (see ACS-4 and SPC-5).

If the URSWRZ bit (see 6.2.2.2) is set to one and the device processes a read command:

- a) with the starting LBA:
  - A) less than the write pointer and the ending LBA greater than or equal to the write pointer in the same zone; or
  - B) greater than or equal to the write pointer and the ending LBA in the same zone;
  - or
- b) with:
  - A) the ending LBA in a different Sequential Write Required zone than the starting LBA; and
  - B) each LBA specified by the command in a Sequential Write Required zone,

then, for any LBA that:

- a) contains user data (see ACS-4) written since the most recent time the Zone Condition was EMPTY for the zone being read, the device shall process the command using the stored data; or
- b) does not contain user data written since the most recent time the Zone Condition was EMPTY for the zone being read, the device shall:
  - A) process the command using the substitute data pattern for Sequential Write Required zones (see 4.6.3.3.3); and
  - B) not write any data for a read command that has the FUA bit, if any, set to one (see ACS-4).

If the URSWRZ bit (see 6.2.2.2) is set to one and the device processes a read command that specifies:

- a) a starting LBA that is in a Sequential Write Required zone; and
- b) an LBA that is not in a Sequential Write Required zone,

then the device shall:

- a) not read any data;
- b) return command completion with an error; and
- c) set the sense key to ILLEGAL REQUEST and the additional sense code to READ BOUNDARY VIOLATION (see ACS-4 and SPC-5).

If the device processes a read command that specifies:

- a) a starting LBA that is not in a Sequential Write Required zone; and
- b) access to an LBA that is in a Sequential Write Required zone,

then the device shall:

- a) not read any data;
- b) return command completion with an error; and
- c) set the sense key to ILLEGAL REQUEST and the additional sense code to ATTEMPT TO READ INVALID DATA (see ACS-4 and SPC-5).

If the device sets the additional sense code to READ BOUNDARY VIOLATION or ATTEMPT TO READ INVALID DATA, then the device shall set the LBA field in the error outputs to the value of the write pointer associated with the Sequential Write Required zone specified by the starting LBA in the command.

#### 4.6.3.3.3 Substitute data pattern for Sequential Write Required zones

For Sequential Write Required zones, the substitute data pattern is data that shall be set by:

- a) the manufacturer; or
- b) the most recent successful:

- A) OVERWRITE EXT command (see ACS-4), if any; or
- B) SCT Write Same command (see ACS-4), if any, with:
  - a) the FUNCTION CODE field set to 0001h, 0002h, 0101h, or 0102h; and
  - b) the WRITE AND SET SUBSTITUTE DATA PATTERN bit (see 4.7.9.2) set to one.

#### 4.6.3.3.4 Opening Sequential Write Required zones

If the number of zones with Zone Condition EXPLICITLY OPENED is equal to the value in the MAXIMUM NUMBER OF OPEN SEQUENTIAL WRITE REQUIRED ZONES field (see 6.2.2.5), then any command (e.g., a write command) that attempts to cause a zone that is not an open zone to become an open zone is completed with an error as described in 4.6.3.4.

#### 4.6.3.4 Zone Condition state machine

##### 4.6.3.4.1 Overview

There is one Zone Condition state machine for each write pointer zone (see 4.6.3). The Zone Condition state machine (see figure 9) controls the operational characteristics of each write pointer zone. This state machine consists of the following states:

- a) ZC1: Empty state (see 4.6.3.4.2);
- b) ZC2: Implicit\_Open state (see 4.6.3.4.3);
- c) ZC3: Explicit\_Open state (see 4.6.3.4.4);
- d) ZC4: Closed state (see 4.6.3.4.5);
- e) ZC5: Full state (see 4.6.3.4.6);
- f) ZC6: Read\_Only state (see 4.6.3.4.7); and
- g) ZC7: Offline state (see 4.6.3.4.8).

Except as specified in 4.7, the device shall maintain the zone state after processing all resets except a power-on reset (see ACS-4).

After a power-on reset, the initial state for each write pointer zone shall be:

- a) the ZC1: Empty state for a zone in which:
  - A) the write pointer is valid and indicates the lowest LBA in the zone; and
  - B) Non-Sequential Write Resources Active is false;
- b) the ZC4: Closed state for a zone in which:
  - A) the write pointer is valid and indicates an LBA that is not the lowest LBA in the zone; or
  - B) Non-Sequential Write Resources Active is true;
- c) the ZC5: Full state for a zone in Zone Condition FULL before the power-on reset;
- d) the ZC6: Read\_Only state for a zone in Zone Condition READ ONLY before the power-on reset; and
- e) the ZC7: Offline state for a zone in Zone Condition OFFLINE before the power-on reset.

On completion of a sanitize operation (see 4.7.6), for each write pointer zone whose Zone Condition is not changed as a result of processing the sanitize operation, that zone's state shall be:

- a) the ZC1: Empty state for a zone in which:
  - A) the write pointer is valid and indicates the lowest LBA in the zone; and
  - B) Non-Sequential Write Resources Active is false;
- b) the ZC4: Closed state for a zone in which:
  - A) the write pointer is valid and indicates an LBA that is not the lowest LBA in the zone; or
  - B) Non-Sequential Write Resources Active is true;
- c) the ZC5: Full state for a zone in Zone Condition FULL before the sanitize operation;
- d) the ZC6: Read\_Only state for a zone in Zone Condition READ ONLY before the sanitize operation; and
- e) the ZC7: Offline state for a zone in Zone Condition OFFLINE before the sanitize operation.

This state machine includes the functions shown in table 6 to define the processing of tasks and provide interfaces for commands that may affect more than one zone as the result of processing one zone at a time.

**Table 6 — Zone Condition state machine functions**

Function	Description	Reference
Manage Open Zone Resources function	Used by this state machine to limit the number of open zones, if required, by closing a zone with a Zone Condition of IMPLICITLY OPENED	4.6.3.4.9
Explicitly Open Zone function	Used by commands (e.g., the OPEN ZONE EXT command (see 5.2.6)) to open a specified zone	4.6.3.4.10
Close Zone function	Used by commands (e.g., the CLOSE ZONE EXT command (see 5.2.4)) to close a specified zone	4.6.3.4.11
Finish Zone function	Used by commands (e.g., the FINISH ZONE EXT command (see 5.2.5)) to modify a specified zone so that its Zone Condition becomes FULL	4.6.3.4.12
Reset Write Pointer function	Used by commands (e.g., the RESET WRITE POINTER EXT command (see 5.2.8)) to modify a specified zone such that its Zone Condition becomes EMPTY	4.6.3.4.13

The characteristics and zone attributes associated with the state of a zone are summarized in table 7.

**Table 7 — Characteristics and attributes associated with zone state**

State	Zone characteristics and attributes			
	Write pointer valid <sup>a</sup>	Open zone resources in use	Accessible for	
			Reads	Writes
ZC1: Empty state	Yes	No	see <sup>b</sup>	see <sup>c</sup>
ZC2: Implicit_Open state	Yes	Yes	Yes	Yes
ZC3: Explicit_Open state	Yes	Yes	Yes	Yes
ZC4: Closed state	Yes	No	Yes	see <sup>c</sup>
ZC5: Full state	No	No	Yes	see <sup>d</sup>
ZC6: Read_Only state	No	No	Yes	No
ZC7: Offline state	No	No	No	No

<sup>a</sup> A valid write pointer (i.e., Yes) indicates a specific LBA in the zone as described in 4.6.3.1. An invalid write pointer (i.e., No) provides no information.

<sup>b</sup> This zone is accessible for reads if:

- the zone type is Sequential Write Preferred (see 4.6.3.2); or
- the zone type is Sequential Write Required (see 4.6.3.3) and the URSWRZ bit is set to one (see 6.2.2.2).

<sup>c</sup> This zone is accessible for writes if the Manage Open Zone Resources function (see 4.6.3.4.9) is able to complete with success and the zone transitions to the ZC2: Implicit\_Open state.

<sup>d</sup> This zone is accessible for writes if the zone type is Sequential Write Preferred (see 4.6.3.2).

This state machine shall maintain the OZR Available state machine variable to indicate the value returned by a Manage Open Zone Resources function (see 4.6.3.4.9).

Figure 9 shows the Zone Condition state machine.

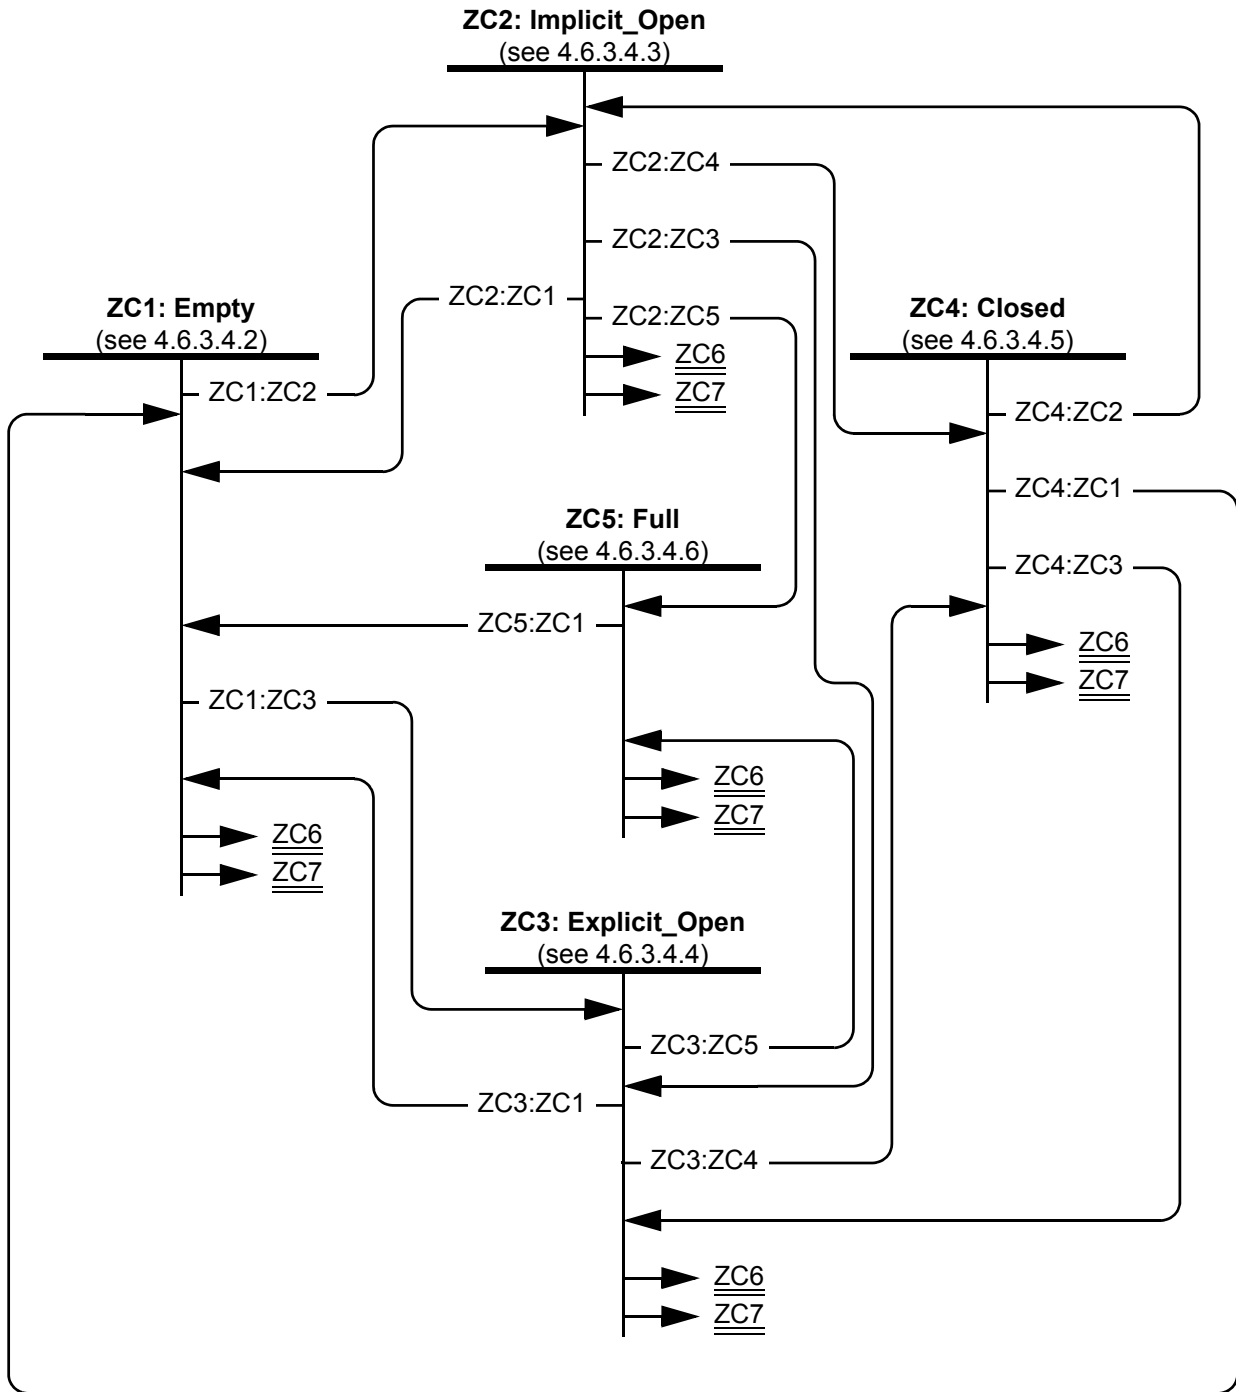



Figure 9 — Zone Condition state machine (part 1 of 2)





<sup>a</sup> The ZC1: Empty state, the ZC2: Implicit\_Open state, the ZC3: Explicit\_Open state, the ZC4: Closed state, and the ZC5: Full state include a transition to this state.

**Figure 9 — Zone Condition state machine (part 2 of 2)**

#### 4.6.3.4.2 ZC1: Empty state

For a zone in this state:

- a) Zone Condition shall be set to EMPTY;
- b) the write pointer is valid and indicates the lowest LBA in the zone;
- c) all zone resources (e.g., the write pointer) shall be persistent zone resources (i.e., open zone resources shall not be used);
- d) RWP Recommended is set as defined in 4.5.5 (i.e., set to false);
- e) Non-Sequential Write Resources Active is set as defined in 4.5.6 (i.e., set to false);
- f) if the device processes:
  - A) a write command with valid parameters;
  - B) an Explicitly Open Zone function (see 4.6.3.4.10); or
  - C) a Finish Zone function (see 4.6.3.4.12),
 then, the device shall:
  - 1) process a Manage Open Zone Resources function (see 4.6.3.4.9) and set the OZR Available state machine variable to the value returned by the function; and
  - 2) if the OZR Available state machine variable is set to FAILURE, the device shall:
    - a) return a failure result for an Explicitly Open Zone function or Finish Zone function; or
    - b) for a write command:
      - A) not write any data;
      - B) return command completion with an error; and
      - C) set the sense key to DATA PROTECT and the additional sense code to INSUFFICIENT ZONE RESOURCES (see ACS-4 and SPC-5);
 and
- g) if the device processes a read command, then the device shall process the read command as described in:
  - A) 4.6.3.2.2 for Sequential Write Preferred zones; or
  - B) 4.6.3.3.2 for Sequential Write Required zones.

**Transition ZC1:ZC2:** The zone shall transition from the ZC1: Empty state to the ZC2: Implicit\_Open state (see 4.6.3.4.3) if:

- a) the device successfully completes:
  - A) a write command; or
  - B) a Finish Zone function (see 4.6.3.4.12);
 and
- b) the OZR Available state machine variable is set to SUCCESS.

**Transition ZC1:ZC3:** The zone shall transition from the ZC1: Empty state to the ZC3: Explicit\_Open state (see 4.6.3.4.4) if:

- a) the device successfully completes an Explicitly Open Zone function (see 4.6.3.4.10); and
- b) the OZR Available state machine variable is set to SUCCESS.

**Transition ZC1:ZC6:** The zone may transition from the ZC1: Empty state to the ZC6: Read\_Only state (see 4.6.3.4.7) for reasons outside the scope of this standard.

**Transition ZC1:ZC7:** The zone may transition from the ZC1: Empty state to the ZC7: Offline state (see 4.6.3.4.8):

- a) as the result of the device detecting a media failure (see 4.7.4); or
- b) for reasons outside the scope of this standard.

#### 4.6.3.4.3 ZC2: Implicit\_Open state

For a zone in this state:

- a) Zone Condition shall be set to IMPLICITLY OPENED;
- b) the write pointer (see 4.6.3.1) is valid;
- c) all zone resources (e.g., the write pointer) shall be open zone resources, persistent zone resources, or both (i.e., this zone uses the open zone resources allocated to it);
- d) RWP Recommended shall be maintained as defined in 4.5.5;
- e) Non-Sequential Write Resources Active shall be maintained as defined in 4.5.6;
- f) if the device processes a write command, then the device shall process the write command as described in:
  - A) 4.6.3.2.1 for Sequential Write Preferred zones; or
  - B) 4.6.3.3.1 for Sequential Write Required zones;
 and
- g) if the device processes a read command, then the device shall process the read command as described in:
  - A) 4.6.3.2.2 for Sequential Write Preferred zones; or
  - B) 4.6.3.3.2 for Sequential Write Required zones.

If this state is entered as the result of a ZC1:ZC2 transition or a ZC4:ZC2 transition, the device shall continue processing the write command or Finish Zone function in this state.

**Transition ZC2:ZC1:** The zone shall transition from the ZC2: Implicit\_Open state to the ZC1: Empty state (see 4.6.3.4.2):

- a) after successful completion of a Reset Write Pointer function (see 4.6.3.4.13) for this zone; or
- b) if the write pointer indicates the lowest LBA in the zone and Non-Sequential Write Resources Active is false, then after successful completion of:
  - A) a Close Zone function (see 4.6.3.4.11) for this zone; or
  - B) a Manage Open Zone Resources function (see 4.6.3.4.9) for a different zone that selects and closes this zone.

**Transition ZC2:ZC3:** The zone shall transition from the ZC2: Implicit\_Open state to the ZC3: Explicit\_Open state (see 4.6.3.4.4) after successful completion of an Explicitly Open Zone function (see 4.6.3.4.10).

**Transition ZC2:ZC4:** If the write pointer does not indicate the lowest LBA in the zone or Non-Sequential Write Resources Active is true, then the zone shall transition from the ZC2: Implicit\_Open state to the ZC4: Closed state (see 4.6.3.4.5) after successful completion of:

- a) a Close Zone function (see 4.6.3.4.11) for this zone; or
- b) a Manage Open Zone Resources function (see 4.6.3.4.9) for a different zone that selects and closes this zone.

**Transition ZC2:ZC5:** For a Sequential Write Required zone, the zone shall transition from the ZC2: Implicit\_Open state to the ZC5: Full state (see 4.6.3.4.6) after successful completion of:

- a) a Finish Zone function (see 4.6.3.4.12); or
- b) a write command with the starting LBA equal to the write pointer that writes the highest LBA in the zone.

For a Sequential Write Preferred zone, the transition from the ZC2: Implicit\_Open state to the ZC5: Full state:

- a) may occur after successful completion of a non-sequential write command in this zone;

- b) shall occur after successful completion of a write command that writes from the write pointer to the highest LBA in this zone; and
- c) shall occur after successful completion of a Finish Zone function in this zone.

**Transition ZC2:ZC6:** The zone may transition from the ZC2: Implicit\_Open state to the ZC6: Read\_Only state (see 4.6.3.4.7) for reasons outside the scope of this standard.

**Transition ZC2:ZC7:** The zone may transition from the ZC2: Implicit\_Open state to the ZC7: Offline state (see 4.6.3.4.8):

- a) as the result of the device detecting a media failure (see 4.7.4); or
- b) for reasons outside the scope of this standard.

#### 4.6.3.4.4 ZC3: Explicit\_Open state

For a zone in this state:

- a) Zone Condition shall be set to EXPLICITLY OPENED;
- b) the write pointer (see 4.6.3.1) is valid;
- c) all zone resources (e.g., the write pointer) shall be open zone resources, persistent zone resources, or both (i.e., this zone uses the open zone resources allocated to it);
- d) RWP Recommended shall be maintained as defined in 4.5.5;
- e) Non-Sequential Write Resources Active shall be maintained as defined in 4.5.6;
- f) if the device processes a write command, then the device shall process the write command as described in:
  - A) 4.6.3.2.1 for Sequential Write Preferred zones; or
  - B) 4.6.3.3.1 for Sequential Write Required zones;
 and
- g) if the device processes a read command, then the device shall process the read command as described in:
  - A) 4.6.3.2.2 for Sequential Write Preferred zones; or
  - B) 4.6.3.3.2 for Sequential Write Required zones.

**Transition ZC3:ZC1:** The zone shall transition from the ZC3: Explicit\_Open state to the ZC1: Empty state (see 4.6.3.4.2):

- a) after successful completion of a Reset Write Pointer function (see 4.6.3.4.13); or
- b) if the write pointer indicates the lowest LBA in the zone and Non-Sequential Write Resources Active is false, then after successful completion of a Close Zone function (see 4.6.3.4.11).

**Transition ZC3:ZC4:** If the write pointer does not indicate the lowest LBA in the zone or Non-Sequential Write Resources Active is true, then the zone shall transition from the ZC2: Implicit\_Open state to the ZC4: Closed state (see 4.6.3.4.5) after successful completion of a Close Zone function (see 4.6.3.4.11).

**Transition ZC3:ZC5:** For a Sequential Write Required zone, the zone shall transition from the ZC3: Explicit\_Open state to the ZC5: Full state (see 4.6.3.4.6) after successful completion of:

- a) a Finish Zone function (see 4.6.3.4.12); or
- b) a write command with the starting LBA equal to the write pointer that writes the highest LBA in the zone.

For a Sequential Write Preferred zone, the transition from the ZC3: Explicit\_Open state to the ZC5: Full state:

- a) may occur after successful completion of a non-sequential write command in this zone;
- b) shall occur after successful completion of a write command that writes from the write pointer to the highest LBA in this zone; and
- c) shall occur after successful completion of a Finish Zone function in this zone.

**Transition ZC3:ZC6:** The zone may transition from the ZC3: Explicit\_Open state to the ZC6: Read\_Only state (see 4.6.3.4.7) for reasons outside the scope of this standard.

**Transition ZC3:ZC7:** The zone may transition from the ZC3: Explicit\_Open state to the ZC7: Offline state (see 4.6.3.4.8):

- a) as the result of the device detecting a media failure (see 4.7.4); or

- b) for reasons outside the scope of this standard.

#### 4.6.3.4.5 ZC4: Closed state

For a zone in this state:

- a) Zone Condition shall be set to CLOSED;
- b) the write pointer (see 4.6.3.1) is valid;
- c) all zone resources (e.g., the write pointer) shall be persistent zone resources (i.e., open zone resources shall not be used);
- d) all user data shall be written to the zone in non-volatile storage (e.g., by performing the equivalent of a flush command (see ACS-4));
- e) RWP Recommended shall be maintained as defined in 4.5.5;
- f) Non-Sequential Write Resources Active shall be maintained as defined in 4.5.6;
- g) if the device processes:
  - A) a write command with valid parameters;
  - B) an Explicitly Open Zone function (see 4.6.3.4.10); or
  - C) a Finish Zone function (see 4.6.3.4.12),
 then, the device shall:
  - 1) process a Manage Open Zone Resources function (see 4.6.3.4.9) and set the OZR Available state machine variable to the value returned by the function; and
  - 2) if the OZR Available state machine variable is set to FAILURE, then the device shall:
    - a) return a failure result for an Explicitly Open Zone function or Finish Zone function; or
    - b) for a write command:
      - A) not write any data;
      - B) return command completion with an error; and
      - C) set the sense key to DATA PROTECT and the additional sense code to INSUFFICIENT ZONE RESOURCES (see ACS-4 and SPC-5);
 and
- h) if the device processes a read command, then the device shall process the read command as described in:
  - A) 4.6.3.2.2 for Sequential Write Preferred zones; or
  - B) 4.6.3.3.2 for Sequential Write Required zones.

**Transition ZC4:ZC1:** The zone shall transition from the ZC4: Closed state to the ZC1: Empty state (see 4.6.3.4.2) after successful completion of a Reset Write Pointer function (see 4.6.3.4.13).

**Transition ZC4:ZC2:** The zone shall transition from the ZC4: Closed state to the ZC2: Implicit\_Open state (see 4.6.3.4.3) if:

- a) the device successfully completes:
  - A) a write command; or
  - B) a Finish Zone function (see 4.6.3.4.12);
 and
- b) the OZR Available state machine variable is set to SUCCESS.

**Transition ZC4:ZC3:** The zone shall transition from the ZC4: Closed state to the ZC3: Explicit\_Open state (see 4.6.3.4.4) if:

- a) the device successfully completes an Explicitly Open Zone function (see 4.6.3.4.10); and
- b) the OZR Available state machine variable is set to SUCCESS.

**Transition ZC4:ZC6:** The zone may transition from the ZC4: Closed state to the ZC6: Read\_Only state (see 4.6.3.4.7) for reasons outside the scope of this standard.

**Transition ZC4:ZC7:** The zone may transition from the ZC4: Closed state to the ZC7: Offline state (see 4.6.3.4.8):

- a) as the result of the device detecting a media failure (see 4.7.4); or
- b) for reasons outside the scope of this standard.

**4.6.3.4.6 ZC5: Full state**

For a zone in this state:

- a) Zone Condition shall be set to FULL;
- b) the write pointer is invalid;
- c) all LBAs in the zone are able to return data in response to a read command;
- d) all zone resources (e.g., the write pointer) shall be persistent zone resources (i.e., open zone resources shall not be used);
- e) RWP Recommended shall be maintained as defined in 4.5.5;
- f) Non-Sequential Write Resources Active shall be maintained as defined in 4.5.6;
- g) if the device processes a write command, then:
  - A) for Sequential Write Preferred zones, the device shall process the write command as described in 4.6.3.2.1; or
  - B) for Sequential Write Required zones, the device shall:
    - a) not write any data;
    - b) return command completion with an error; and
    - c) set the sense key to ILLEGAL REQUEST and the additional sense code to INVALID FIELD IN CDB (see ACS-4 and SPC-5);
- and
- h) if the device processes a read command, then the device shall process the read command as described in:
  - A) 4.6.3.2.2 for Sequential Write Preferred zones; or
  - B) 4.6.3.3.2 for Sequential Write Required zones.

**Transition ZC5:ZC1:** The zone shall transition from the ZC5: Full state to the ZC1: Empty state (see 4.6.3.4.2) after successful completion of a Reset Write Pointer function (see 4.6.3.4.13).

**Transition ZC5:ZC6:** The zone may transition from the ZC5: Full state to the ZC6: Read\_Only state (see 4.6.3.4.7) for reasons outside the scope of this standard.

**Transition ZC5:ZC7:** The zone may transition from the ZC5: Full state to the ZC7: Offline state (see 4.6.3.4.8):

- a) as the result of the device detecting a media failure (see 4.7.4); or
- b) for reasons outside the scope of this standard.

**4.6.3.4.7 ZC6: Read\_Only state**

For a zone in this state:

- a) Zone Condition shall be set to READ ONLY;
- b) the write pointer is invalid;
- c) all zone resources (e.g., the write pointer) shall be persistent zone resources (i.e., open zone resources shall not be used);
- d) RWP Recommended is set as defined in 4.5.5 (i.e., set to false);
- e) Non-Sequential Write Resources Active is set as defined in 4.5.6 (i.e., set to false);
- f) if the device processes a write command or a FINISH ZONE EXT command (see 5.2.5), the device shall:
  - A) not write any data;
  - B) return command completion with an error; and
  - C) set the sense key to DATA PROTECT and the additional sense code to ZONE IS READ ONLY (see ACS-4 and SPC-5);
- and
- g) if the device processes a read command, then the device shall process the read command as described in:
  - A) 4.6.3.2.2 for Sequential Write Preferred zones; or
  - B) 4.6.3.3.2 for Sequential Write Required zones.

With the exception of transitions to the ZC7: Offline state (see 4.6.3.4.8), transitions out of this state are outside the scope of this standard.

**Transition ZC6:ZC7:** The zone may transition from the ZC6: Read\_Only state to the ZC7: Offline state (see 4.6.3.4.8):

- a) as the result of the device detecting a media failure (see 4.7.4);
- b) as the result of performing a sanitize operation (see 4.7.6); or
- c) for reasons outside the scope of this standard.

#### 4.6.3.4.8 ZC7: Offline state

For a zone in this state:

- a) Zone Condition shall be set to OFFLINE;
- b) the write pointer is invalid;
- c) all zone resources (e.g., the write pointer) shall be persistent zone resources (i.e., open zone resources shall not be used);
- d) RWP Recommended is set as defined in 4.5.5 (i.e., set to false);
- e) Non-Sequential Write Resources Active is set as defined in 4.5.6 (i.e., set to false); and
- f) if a device processes a read command, a write command, or a FINISH ZONE EXT command (see 5.2.5), then the device shall:
  - A) return command completion with an error; and
  - B) set the sense key to MEDIUM ERROR and the additional sense code to ZONE IS OFFLINE (see ACS-4 and SPC-5).

Transitions out of this state are outside the scope of this standard.

#### 4.6.3.4.9 Manage Open Zone Resources function

##### 4.6.3.4.9.1 Overview

The Manage Open Zone Resources function provides a way for the Zone Condition state machine to allow a zone to become an open zone while maintaining the limits on the number of open zones that the device indicates in:

- a) the OPTIMAL NUMBER OF OPEN SEQUENTIAL WRITE PREFERRED ZONES field (see 6.2.2.3); and
- b) the MAXIMUM NUMBER OF OPEN SEQUENTIAL WRITE REQUIRED ZONES field (see 6.2.2.5).

The processing of a Manage Open Zone Resources function is requested only within the Zone Condition state machine.

The Manage Open Zone Resources function shall return:

- a) SUCCESS to indicate that sufficient open zone resources are available to allow a Zone Condition to become IMPLICITLY OPENED or EXPLICITLY OPENED; or
- b) FAILURE to indicate that insufficient open zone resources are available to allow a Zone Condition to become IMPLICITLY OPENED or EXPLICITLY OPENED.

The Manage Open Zone Resources function:

- a) determines whether open zone resources are available to allow a zone to become an open zone; and
- b) may close a zone with a Zone Condition of IMPLICITLY OPENED in order to make open zone resources available.

As detailed in 4.6.3.4.9, the Manage Open Zone Resources function returns SUCCESS if:

- a) closing a zone is not required in order to limit the number of open zones; or
- b) a zone is closed in order to limit the number of open zones.

Otherwise, the Manage Open Zone Resources function returns FAILURE.

##### 4.6.3.4.9.2 Processing a sequential write preferred zone

For a sequential write preferred zone:

- a) if equation (1) evaluates to true, the device may select a sequential write preferred zone with a Zone Condition of IMPLICITLY OPENED and process a Close Zone function (see 4.6.3.4.11) on that zone; and
- b) the function shall return SUCCESS.

$$(\text{optimal} \leq (\text{explicit}_p + \text{implicit}_p)) \text{ AND } (0 < \text{implicit}_p) \quad (1)$$

where:

$\text{explicit}_p$	is the number of Sequential Write Preferred zones (see 4.6.3.2) with a Zone Condition of EXPLICITLY OPENED;
$\text{implicit}_p$	is the number of Sequential Write Preferred zones with a Zone Condition of IMPLICITLY OPENED; and
$\text{optimal}$	is the contents of the OPTIMAL NUMBER OF OPEN SEQUENTIAL WRITE PREFERRED ZONES field (see 6.2.2.3).

#### 4.6.3.4.9.3 Processing a sequential write required zone

If equation (2) evaluates to true, then:

- a) the device should not select a zone to be closed; and
- b) the function shall return SUCCESS.

$$\text{max} > (\text{explicit}_r + \text{implicit}_r) \quad (2)$$

where:

$\text{explicit}_r$	is the number of Sequential Write Required zones (see 4.6.3.3) with a Zone Condition of EXPLICITLY OPENED;
$\text{implicit}_r$	is the number of Sequential Write Required zones with a Zone Condition of IMPLICITLY OPENED; and
$\text{max}$	is the contents of the MAXIMUM NUMBER OF OPEN SEQUENTIAL WRITE REQUIRED ZONES field (see 6.2.2.5).

If equation (3) evaluates to true, the function shall return FAILURE.

$$\text{max} \leq \text{explicit}_r \quad (3)$$

where:

$\text{explicit}_r$	is the number of Sequential Write Required zones (see 4.6.3.3) with a Zone Condition of EXPLICITLY OPENED; and
$\text{max}$	is the contents of the MAXIMUM NUMBER OF OPEN SEQUENTIAL WRITE REQUIRED ZONES field (see 6.2.2.5).

If equation (4) evaluates to true, then:

- 1) the device shall choose a zone with a Zone Condition of IMPLICITLY OPENED and process a Close Zone function (see 4.6.3.4.11) on that zone; and
- 2) the function shall return SUCCESS.

$$(\text{max} \leq (\text{explicit}_r + \text{implicit}_r)) \text{ AND } (0 < \text{implicit}_r) \quad (4)$$

where:

$\text{explicit}_r$	is the number of Sequential Write Required zones (see 4.6.3.3) with a Zone Condition of EXPLICITLY OPENED;
$\text{implicit}_r$	is the number of Sequential Write Required zones with a Zone Condition of IMPLICITLY OPENED; and
$\text{max}$	is the contents of the MAXIMUM NUMBER OF OPEN SEQUENTIAL WRITE REQUIRED ZONES field (see 6.2.2.5).

#### 4.6.3.4.10 Explicitly Open Zone function

The Explicitly Open Zone function prepares a write pointer zone for writing (e.g., by allocating any open zone resources associated with the write pointer zone).

For the specified zone:

- 1) the device shall process the Manage Open Zone Resources function (see 4.6.3.4.9);
- 2) if the Manage Open Zone Resources function returned FAILURE, then the Explicitly Open Zone function shall return FAILURE; and

- 3) if the Manage Open Zone Resources function returned SUCCESS, then:
  - 1) the Zone Condition for that zone shall become EXPLICITLY OPENED; and
  - 2) the Explicitly Open Zone function shall return SUCCESS.

An Explicitly Open Zone function may be processed on a zone with a Zone Condition of EXPLICITLY OPENED. This has no effect and is not considered to be an error.

#### 4.6.3.4.11 Close Zone function

A successful Close Zone function annuls any preparations made for writing an open zone (e.g., by deallocating any open zone resources associated with the write pointer zone).

As a result of a successful Close Zone function, the following become non-volatile:

- a) the write pointer for the specified write pointer zone; and
- b) all user data stored in the specified write pointer zone.

For the specified zone, the Zone Condition state machine processing of this function (e.g., as shown in the ZC2: Implicit\_Open state (see 4.6.3.4.3)) results in the Zone Condition for the specified zone becoming:

- a) EMPTY, if the write pointer indicates the lowest LBA in the zone and Non-Sequential Write Resources Active is false; or
- b) CLOSED, if the write pointer does not indicate the lowest LBA in the zone or Non-Sequential Write Resources Active is true.

A Close Zone function may be processed on a zone with a Zone Condition of CLOSED. Such an operation has no effect and is not considered to be an error.

#### 4.6.3.4.12 Finish Zone function

The result of a successful Finish Zone function is:

- a) the Zone Condition becomes FULL; and
- b) all user data in the zone becomes non-volatile and available for reading.

The device shall process the Finish Zone function for the specified zone as follows:

- 1) if the Zone Condition is EMPTY or CLOSED, perform the processing described in 4.6.3.4.2 and 4.6.3.4.5; and
- 2) if the Zone Condition after processing step 1) is EXPLICITLY OPENED or IMPLICITLY OPENED, then:
  - A) subsequent reads that specify logical sectors in the zone that have not been written since the last time the Zone Condition was EMPTY shall return the substitute data pattern (see 4.6.3.2.3 and 4.6.3.3.3); and
  - B) the Zone Condition for the specified zone becomes FULL.

A finish zone function that is performed on a zone with a Zone Condition of FULL has no effect and is not considered to be an error.

#### 4.6.3.4.13 Reset Write Pointer function

A successful Reset Write Pointer function results in the Zone Condition becoming EMPTY (i.e., the write pointer is set to the lowest LBA in the zone and the write pointer becomes valid) for the specified zone.

A Reset Write Pointer function that is performed on a zone with a Zone Condition of EMPTY has no effect and is not considered to be an error.

## 4.7 Additional features for the Host Aware Zones feature set and the Host Managed Zones feature set

### 4.7.1 Introduction

Support for the Host Aware Zones feature set (see 4.3) and the Host Managed Zones feature set (see 4.4) modifies features defined in ACS-4.



#### 4.7.2 Zoned device internal resource management

A device that supports the Host Aware Zones feature set (see 4.3) or the Host Managed Zones feature set (see 4.4) requires internal resources (e.g., persistent zone resources) to maintain each zone. Insufficient resources may result in degraded functionality (e.g., reduced performance, increased power consumption, or increased reporting of write errors).

A RESET bit set to one in a zone descriptor (see 5.2.7.6.6) indicates that an insufficient resources condition has been detected for that zone.

A host that detects the RESET bit being set to one may respond by sending a RESET WRITE POINTER EXT command (see 5.2.8) that specifies the affected zone. Before sending the RESET WRITE POINTER EXT command, the host may copy the user data in the affected zone to another location.

#### 4.7.3 Unexpected power removal

If power is removed from a zoned device and:

- a) there are partially completed commands that result in writing the non-volatile media;
- b) there is user data in volatile write cache; or
- c) zone attributes for completed write commands are not stored in persistent zone resources,

then an unexpected power removal condition has occurred.

An unexpected power removal condition may result in zone attributes for partially completed write commands being updated to reflect the partial user data written to the medium (e.g., a write pointer may indicate the LBA plus one of the last logical sector that the device attempted to write before the unexpected power removal condition).

If:

- 1) the device has non-volatile media;
- 2) all write cache is non-volatile (i.e., the NON-VOLATILE WRITE CACHE ENABLED bit (see ACS-4) is set to one) or the volatile write cache is disabled (i.e., the VOLATILE WRITE CACHE ENABLED bit (see ACS-4) is cleared to zero);
- 3) a write command completes without error; and
- 4) an unexpected power removal occurs,

then, after power is restored, a read command that completes without error for any LBA included in that write command shall return the user data that was written by that write command before the unexpected power removal.

If:

- 1) the device has non-volatile media;
- 2) a write FUA command completes without error; and
- 3) an unexpected power removal occurs,

then, after power is restored, a read command that completes without error for any LBA included in that write FUA command shall return the user data that was written by that write FUA command before the unexpected power removal.

If:

- 1) the device has non-volatile media;
- 2) a write command completes without error;
- 3) a subsequent flush command completes without error; and
- 4) an unexpected power removal occurs,

then, after power is restored, a read command that completes without error for any LBA included in that write command shall return the user data that was written by that write command before the unexpected power removal.

For Sequential Write Required zones (see 4.6.3.3), if an unexpected power removal results in unwritten logical sectors in LBAs lower than the write pointer for a zone and read commands access those unwritten logical sectors, then the device shall:

- a) return command completion with an error; and
- b) set the sense key to MEDIUM ERROR and the additional sense code to UNWRITTEN DATA IN ZONE (see ACS-4 and SPC-5).

After an unexpected power removal condition, the data transferred to the host by a REPORT ZONES EXT command (see 5.2.7) may contain the write pointers that do not correspond to the ending LBA plus one of the last command to report completion prior to the unexpected power removal condition.

#### 4.7.4 Media failure

If a zoned device detects a media failure that affects one or more write pointer zones, the device may indicate that those zones are unavailable for reading or writing by changing the Zone Condition to OFFLINE.

If the device processes a write command that specifies a write pointer zone that results in the additional sense code being set to (see ACS-4 and SPC-5):

- a) WRITE ERROR - RECOVERY NEEDED;
- b) WRITE ERROR - RECOVERY SCAN NEEDED; or
- c) WRITE ERROR - INSUFFICIENT ZONE RESOURCES,

then the device shall set the LBA field in the error outputs to the lowest numbered LBA where a media failure may have occurred during processing of that write command.

If the additional sense code is set to WRITE ERROR - RECOVERY SCAN NEEDED, the host should:

- a) read all the logical blocks between the LBA indicated by the LBA field in the error outputs and the LBA indicated by the write pointer minus one; and
- b) take recovery actions for the LBAs, if any, that return errors.

#### 4.7.5 Interactions with volatile caches

While processing write commands there is a period of time during which the user data may be lost as a result of using volatile write cache (see ACS-4) if:

- a) an unexpected power removal occurs (see 4.7.3); or
- b) a hardware failure occurs.

If an error occurs while the device is writing to the medium in a Sequential Write Required zone (see 4.6.3.3) and that error is reported as a deferred error (see ACS-4), then the device may invalidate cached user data for LBAs in that zone that are greater than the LBA reported with the deferred error. This invalidation may occur for data cached in both volatile and non-volatile caches.

The host may affect behavior of the volatile write cache with:

- a) the CLOSE ZONE command (see 5.2.4);
- b) the FINISH ZONE command (see 5.2.5);
- c) read FUA commands;
- d) write FUA commands; and
- e) flush commands.

If all write cache is non-volatile (i.e., the NON-VOLATILE WRITE CACHE bit (see ACS-4) is set to one), then all zone resources shall be persistent zone resources.

If volatile write cache is disabled (i.e., the VOLATILE WRITE CACHE ENABLED bit (see ACS-4) is cleared to zero), then all zone resources shall be persistent zone resources.

If volatile write cache is enabled (i.e., the VOLATILE WRITE CACHE ENABLED bit is set to one) and the device processes:

- a) a CLOSE ZONE EXT command with the CLOSE ALL bit set to one (see 5.2.4.3.2);
- b) a FINISH ZONE EXT command with the FINISH ALL bit set to one (see 5.2.5.3.2);
- c) a RESET WRITE POINTER EXT command with the RESET ALL bit set to one (see 5.2.8.3.2);

- d) a flush command; or
- e) a STANDBY IMMEDIATE command (see ACS-4),

then, for all open write pointer zones, all zone resources shall become persistent zone resources before returning command completion without error.

If volatile write cache is enabled, then prior to a transition to a power condition that prevents access to the media (e.g., before a hard drive stops its spindle motor during a transition to the Standby\_z power condition (see ACS-4)) for all open write pointer zones all zone resources shall become persistent zone resources.

If volatile write cache is enabled (i.e., the VOLATILE WRITE CACHE ENABLED bit is set to one) and the device processes:

- a) a CLOSE ZONE EXT command with the CLOSE ALL bit cleared to zero (see 5.2.4.3.2);
- b) a FINISH ZONE EXT command with the FINISH ALL bit cleared to zero (see 5.2.5.3.2);
- c) a RESET WRITE POINTER EXT command with the RESET ALL bit cleared to zero (see 5.2.8.3.2);
- d) a read FUA command;
- e) a write FUA command; or
- f) a write stream command with the FLUSH bit set to one (see ACS-4),

then, for the affected open write pointer zones, all zone resources shall become persistent zone resources before returning command completion without error.

If:

- a) the write-read-verify feature set (see ACS-4) is enabled; and
- b) the processing of a write command requires the device to read from the media before returning command completion without error,

then, for the affected open write pointer zones, all zone resources shall become persistent zone resources before returning command completion without error for that write command.

If volatile write cache is enabled and the device returns command completion without error for an IDLE IMMEDIATE command (see ACS-4), then all open zone resources that are volatile may remain volatile.

#### 4.7.6 Interactions with the Sanitize Device feature set

If a device that supports the Host Aware Zones feature set (see 4.3) or the Host Managed Zones feature set (see 4.4) also supports the Sanitize Device feature set (see ACS-4), then the additional requirements described in this subclause apply.

The ZONED NO RESET bit shall be supported as described in ACS-4 for each of the following commands that are supported:

- a) the CRYPTO SCRAMBLE EXT command (see ACS-4);
- b) the BLOCK ERASE EXT (see ACS-4) command; and
- c) the OVERWRITE EXT command (see ACS-4).

A CRYPTO SCRAMBLE EXT command, BLOCK ERASE EXT command, or OVERWRITE EXT command affects all zones as follows:

- a) the specified sanitize operation is performed as specified in ACS-4 for each Conventional zone (see 4.6.2);
- b) the specified sanitize operation is performed as specified in ACS-4 for each write pointer zone (see 4.6.3) and shall include processing of the ZONED NO RESET bit as described in ACS-4 if a sanitize operation:
  - A) is successful; or
  - B) fails and:
    - a) the Failure Mode Policy value (see ACS-4) allows successful processing of a SANITIZE STATUS EXT command with the CLEAR SANITIZE OPERATION FAILED bit set to one; and
    - b) a failed sanitize operation is followed by a SANITIZE STATUS EXT command with the CLEAR SANITIZE OPERATION FAILED bit set to one;

- c) if the Zone Condition is READ ONLY before the specified sanitize operation is performed, the Zone Condition shall not change to any value except OFFLINE as a result of an error in the specified sanitize operation; and
- d) if the specified sanitize operation completed without error, the zone's state shall be set as described in 4.6.3.4.1 with the following exceptions:
  - A) the Zone Condition shall become OFFLINE for every write pointer zone that was not able to be sanitized (e.g., due to an inability to write to that zone); and
  - B) for reasons outside the scope of this standard, the Zone Condition may become OFFLINE for a write pointer zone for which the Zone Condition was READ ONLY before that command was processed.

An OVERWRITE EXT command that completes without an error modifies the substitute data pattern (see 4.6.3.2.3 and 4.6.3.3.3).

#### 4.7.7 Interactions with the Power Management feature set

As the result of a Power Management Transition PM3:PM2 (see ACS-4) the state of each write pointer zone may be set as indicated for the initial state after a power-on reset (see 4.6.3.4.1).

#### 4.7.8 Interactions with the SECURITY ERASE UNIT command

If a device supports the Host Aware feature set (see 4.3) and the device processes a SECURITY ERASE UNIT command (see ACS-4), then prior to returning command completion without error:

- a) the device shall set Non-Sequential Write Resources Active to false for all write pointer zones (see 4.6.3); and
- b) if the ZAC SECURITY OPTION bit (see ACS-4) is:
  - A) cleared to zero, then the device shall cause the Zone Condition to become EMPTY for all write pointer zones; or
  - B) set to one, then the device shall cause the Zone Condition to become FULL for all write pointer zones.

#### 4.7.9 Interactions with the SCT Write Same command

##### 4.7.9.1 Overview

If the FUNCTION CODE field is set to 0001h, 0002h, 0101h, or 0102h in an SCT Write Same command (see ACS-4), then the ZAC OPTIONS field (see table 8) adds zoned device support capabilities to the processing of the command. Unless otherwise specified, specifying a zoned device capability using the ZAC OPTIONS field shall not affect the processing of an SCT Write Same command as described in ACS-4.

**Table 8 — ZAC OPTIONS field**

Bit	Meaning
15	WRITE AND SET SUBSTITUTE DATA PATTERN bit (see 4.7.9.2)
14:4	Reserved
3:0	ZONE POST PROCESSING field (see 4.7.9.3)

If the FUNCTION CODE field is not set to 0001h, 0002h, 0101h, or 0102h and the ZAC OPTIONS field is not cleared to zero, then the device should:

- a) complete the command with an error; or
- b) process the command as if the ZAC OPTIONS field is cleared to zero.

##### 4.7.9.2 WRITE AND SET SUBSTITUTE DATA PATTERN bit

A zoned device shall process an SCT Write Same command as follows:

- 1) perform the actions specified for the SCT Write Same command (see ACS-4);
- 2) if the WRITE AND SET SUBSTITUTE DATA PATTERN bit is set to one, set the substitute data pattern; and
- 3) process the ZONE POST PROCESSING field (see 4.7.9.3).

If an error occurs during any of these steps, the subsequent steps shall not be performed.

**4.7.9.3 ZONE POST PROCESSING field**

After successfully processing the SCT Write Same command as described in ACS-4, the ZONE POST PROCESSING field (see table 9) specifies additional processing that shall be done by the device.

**Table 9 — ZONE POST PROCESSING field**

Code	Description
0h	<b>None:</b> The device shall process the SCT Write Same command as described in ACS-4 with no additional processing.
1h	<b>Close Affected Zones:</b> The device shall process a Close Zone function (see 4.6.3.4.11) for each zone to which user data was written during the processing of the SCT Write Same command.
2h	<b>Reset Affected Zones:</b> The device shall process a Reset Write Pointer function (see 4.6.3.4.13) for each zone to which user data was written during the processing of the SCT Write Same command.
3h to 8h	Reserved
9h	<b>Close All Zones:</b> The device shall process a Close Zone function (see 4.6.3.4.11) for each zone with a Zone Condition of IMPLICITLY OPENED or EXPLICITLY OPENED.
Ah	<b>Reset All Zones:</b> The device shall process a Reset Write Pointer function (see 4.6.3.4.13) for each zone with a Zone Condition of IMPLICITLY OPENED, EXPLICITLY OPENED, CLOSED, or FULL.
Bh..Fh	Reserved

## 5 Command descriptions

### 5.1 Command description introduction

#### 5.1.1 Overview

Clause 5 describes the commands defined by this standard as part of the Host Aware Zones feature set (see 4.3) and the Host Managed Zones feature set (see 4.4).

Unless otherwise defined, all aspects of the command descriptions in clause 5 are as defined in ACS-4.

ATA commands are delivered using the following fields (see ACS-4):

- a) FEATURE;
- b) COUNT;
- c) LBA;
- d) DEVICE; and
- e) COMMAND.

ATA commands may use the ICC field and the AUXILIARY field (see ACS-4).

Field lengths are different based on the type of command (see ACS-4).

This standard describes the ATA command set in a transport independent fashion. Each command is defined by a series of subclauses as described in 5.1.2 through 5.1.5.

#### 5.1.2 Command Name – Command Code [/Subcommand Code], Command Protocol

The heading for each command starts with the name of the command. The name is followed by “–” and then the command code, subcommand code if applicable, and protocol used to process the command.

EXAMPLE - A heading reads:

READ SECTOR(S) – 20h, PIO Data-In

In this example heading the name of the command is READ SECTOR(S). The command code is 20h. The protocol used to transfer the data is PIO Data-In.

Protocols are defined in ATA8-AAM. The transport protocol standards define the implementation of each protocol.

#### 5.1.3 Inputs

The Inputs subclause contains a table showing the inputs for the command. An example command structure is shown in table 10.

**Table 10 — Example Command Structure**

Field	Description
FEATURE	Each transport standard defines how the FEATURE field is mapped for proper functionality. Each transport standard also defines how 28-bit commands are mapped differently than 48-bit commands.
COUNT	Each transport standard defines how the COUNT field is mapped for proper functionality. Each transport standard also defines how 28-bit commands are mapped differently than 48-bit commands.
LBA	For many commands, the LBA field contains the LBA of the first logical sector to be transferred. Each transport standard defines how the LBA field is mapped to the appropriate fields or registers.
DEVICE	Each transport standard defines how the DEVICE field bits 7:4 are mapped. Bits 3:0 are marked reserved in every reference to the DEVICE field.
COMMAND	The COMMAND field contains the command code.

5.1.4 Normal Outputs

This is an example Normal Output. A command with Normal Outputs does not return command completion with an error. Therefore, the ERROR field in the Normal Outputs is reserved in every command. The COUNT field and LBA field may be reserved. However, in some commands the COUNT field and LBA field may have return parameters in Normal Outputs. The STATUS field shows the DEVICE FAULT bit, the ALIGNMENT ERROR bit, the SENSE DATA AVAILABLE bit, and the ERROR bit. Bit 7, bit 6, and bit 3 of the STATUS field are marked Transport Dependent in many of the Normal Outputs.

Table 11 — Example Normal Output

Field	Description
ERROR	Reserved
COUNT	Reserved
LBA	Reserved
DEVICE	<p><b>Bit Description</b></p> <ul style="list-style-type: none"> <li>7 Obsolete</li> <li>6 N/A</li> <li>5 Obsolete</li> <li>4 Transport Dependent – See ACS-4</li> <li>3:0 Reserved</li> </ul>
STATUS	<p><b>Bit Description</b></p> <ul style="list-style-type: none"> <li>7:6 Transport Dependent – See ACS-4</li> <li>5 DEVICE FAULT bit – See ACS-4</li> <li>4 N/A</li> <li>3 Transport Dependent – See ACS-4</li> <li>2 N/A or ALIGNMENT ERROR bit – See ACS-4</li> <li>1 SENSE DATA AVAILABLE bit – See ACS-4</li> <li>0 ERROR bit – See ACS-4</li> </ul>

5.1.5 Error Outputs

The Error Outputs subclause shows the ERROR field, COUNT field, LBA field, and STATUS field. An Error Output occurs when any bit in the STATUS field (e.g., the ERROR bit, or the DEVICE FAULT bit) is set to one, indicating that an error occurred. If the ERROR bit is set to one, the ERROR field indicates the type of error that occurred.

Table 12 — Example Error Output

Field	Description
ERROR	<p><b>Bit Description</b></p> <ul style="list-style-type: none"> <li>7 INTERFACE CRC bit – See ACS-4</li> <li>6 UNCORRECTABLE ERROR bit – See ACS-4</li> <li>5 Obsolete</li> <li>4 ID NOT FOUND bit – See ACS-4</li> <li>3 Obsolete</li> <li>2 ABORT bit – See ACS-4</li> <li>1 Obsolete</li> <li>0 Obsolete</li> </ul>
COUNT	Reserved
LBA	LBA of First Unrecoverable Error – See ACS-4
DEVICE	<p><b>Bit Description</b></p> <ul style="list-style-type: none"> <li>7 Obsolete</li> <li>6 N/A</li> <li>5 Obsolete</li> <li>4 Transport Dependent – See ACS-4</li> <li>3:0 Reserved</li> </ul>
STATUS	<p><b>Bit Description</b></p> <ul style="list-style-type: none"> <li>7:6 Transport Dependent – See ACS-4</li> <li>5 DEVICE FAULT bit – See ACS-4</li> <li>4 N/A</li> <li>3 Transport Dependent – See ACS-4</li> <li>2:1 N/A</li> <li>0 ERROR bit – See ACS-4</li> </ul>



## 5.2 Zone management commands

### 5.2.1 Overview

Each zone management command (see table 13) defines a mandatory or optional management function to be processed by a zoned device using:

- a) a command template using the COMMAND field input; and
- b) a command protocol (see ACS-4).

**Table 13 — Zone management commands**

Command	Support	Template	Command protocol	Reference
CLOSE ZONE EXT	Mandatory	Out	Non-Data	5.2.4
FINISH ZONE EXT	Mandatory	Out	Non-Data	5.2.5
OPEN ZONE EXT	Mandatory	Out	Non-Data	5.2.6
REPORT ZONES EXT	Mandatory	In	DMA	5.2.7
RESET WRITE POINTER EXT	Mandatory	Out	Non-Data	5.2.8
Key: In = ZAC Management In (see 5.2.2) Out = ZAC Management Out (see 5.2.3)				

## 5.2.2 ZAC Management In command template, 4Ah, DMA

### 5.2.2.1 Feature Set

This 48-bit command is for devices that support the Host Aware Zones feature set (see 4.3) or the Host Managed Zones feature set (see 4.4).

### 5.2.2.2 Description

Commands based on the ZAC Management In command template transfer one or more 512-byte pages of data from the device to the host. The data to be transferred is specified by the contents of the ZM\_ACTION field (see 5.2.2.3) and other command inputs.

The device processes any command that uses the ZAC Management In command template in the NCQ feature set environment (see ACS-4) if that command is encapsulated in a RECEIVE FPDMA QUEUED command (see ACS-4) with the inputs encapsulated as shown in 5.2.2.5.

### 5.2.2.3 Inputs

#### 5.2.2.3.1 Overview

Table 14 shows the command inputs defined for all commands that use the ZAC Management In command template.

**Table 14 — ZAC Management In command template inputs**

Field	Description
FEATURE	<p><b>Bit Description</b></p> <p>15:8 Action specific – See 5.2.2.3.2</p> <p>7:5 Reserved</p> <p>4:0 ZM_ACTION field – See 5.2.2.3.3</p>
COUNT	RETURN PAGE COUNT field – See 5.2.2.3.4
LBA	Action specific – See 5.2.2.3.2
DEVICE	<p><b>Bit Description</b></p> <p>7 Obsolete</p> <p>6 Shall be set to one</p> <p>5 Obsolete</p> <p>4 Transport Dependent – See ACS-4</p> <p>3:0 Reserved</p>
COMMAND	7:0 4Ah

#### 5.2.2.3.2 Action specific fields

Each zone management command (see 5.2.1) defines the usage, if any, of the action specific fields in that command's inputs.

**5.2.2.3.3 ZM\_ACTION field**

The ZM\_ACTION field (see table 15) and the COMMAND field specify:

- a) the zone management command that the device shall process;
- b) the definitions for action specific fields (see 5.2.2.3.2) in that command's input, if any; and
- c) the definitions of the normal outputs and error outputs for that command.

**Table 15 — ZM\_ACTION field**

Code	Command	Reference
00h	REPORT ZONES EXT	5.2.7
01h	CLOSE ZONE EXT	5.2.4
02h	FINISH ZONE EXT	5.2.5
03h	OPEN ZONE EXT	5.2.6
04h	RESET WRITE POINTER EXT	5.2.8
all others	Reserved	
Key:		
In = ZAC Management In (see 5.2.2)		
Out = ZAC Management Out (see 5.2.3)		

**5.2.2.3.4 RETURN PAGE COUNT field**

The device shall return the number of 512-byte pages specified in the RETURN PAGE COUNT field. Pad bytes are appended as needed to meet this requirement. Pad bytes shall have a value of 00h.

Unless otherwise specified, if the RETURN PAGE COUNT field specifies fewer 512-byte pages than the device has available to be returned, the device:

- a) shall truncate the returned data to the specified number of 512-byte pages; and
- b) shall not modify any of the returned data as a result of the truncation.

If the RETURN PAGE COUNT field specifies more 512-byte pages than the device has available to be returned, then 512 bytes of all zeros shall be returned for those pages.

The value 0000h is reserved in the RETURN PAGE COUNT field.

**5.2.2.4 Normal Outputs and Error Outputs**

The normal outputs and error outputs for a command that uses the ZAC Management In command template are defined based on the contents of the ZM\_ACTION field (see 5.2.2.3).

**5.2.2.5 NCQ encapsulation for a DMA ZAC Management In command**

If a command that uses the ZAC Management In command template is processed in an NCQ environment as subcommand 02h of RECEIVE FPDMA QUEUED command (see ACS-4), then the command inputs are encapsulated as shown in table 16.

**Table 16 — RECEIVE FPDMA QUEUED command encapsulation for the inputs used by the ZAC Management In command**

RECEIVE FPDMA QUEUED field (see ACS-4)	ZAC Management In field (see table 14), if any
LBA	LBA
AUXILIARY (31:16)	Reserved
AUXILIARY (15:0)	FEATURE (15:0)

### 5.2.3 ZAC Management Out command template, 9Fh, Non-Data or DMA

#### 5.2.3.1 Feature Set

This 48-bit command is for devices that support the Host Aware Zones feature set (see 4.3) or the Host Managed Zones feature set (see 4.4).

#### 5.2.3.2 Description

Commands based on the ZAC Management Out command template transfer zero or more 512-byte pages of data from the host to the device. What data is transferred is specified by the contents of the ZM\_ACTION field (see 5.2.2.3) and other command inputs.

The device processes any command that uses the ZAC Management Out command template in the NCQ feature set environment (see ACS-4) if that:

- a) Non-Data command is encapsulated in a NCQ NON-DATA command (see ACS-4) with the inputs encapsulated as shown in 5.2.3.5; or
- b) DMA command is encapsulated in a SEND FPDMA QUEUED command (see ACS-4) with the inputs encapsulated as shown in 5.2.3.6.

#### 5.2.3.3 Inputs

##### 5.2.3.3.1 Overview

Table 17 shows the command inputs defined for all commands that use the ZAC Management Out, Non-Data command template or the ZAC Management Out, DMA command template.

**Table 17 — ZAC Management Out command template inputs**

Field	Description
FEATURE	<p><b>Bit Description</b></p> <p>15:8 Action specific – See 5.2.2.3.2</p> <p>7:5 Reserved</p> <p>4:0 ZM_ACTION field – See 5.2.2.3.3</p>
COUNT	<p>ZAC Management Out, Non-Data (see table 13): Action specific – See 5.2.2.3.2</p> <p>ZAC Management Out, DMA (see table 13): PAGES TO SEND field – See 5.2.3.3.2</p>
LBA	Action specific – See 5.2.2.3.2
DEVICE	<p><b>Bit Description</b></p> <p>7 Obsolete</p> <p>6 Shall be set to one</p> <p>5 Obsolete</p> <p>4 Transport Dependent – See ACS-4</p> <p>3:0 Reserved</p>
COMMAND	7:0 9Fh

##### 5.2.3.3.2 PAGES TO SEND field

The PAGES TO SEND field specifies the number of 512-byte pages that the device shall transfer from the host.

If the PAGES TO SEND field is cleared to zero in a ZAC Management Out, DMA command, the device shall return command aborted.

### 5.2.3.4 Normal Outputs and Error Outputs

The normal outputs and error outputs for a command that uses the ZAC Management Out command template are defined based on the contents of the ZM\_ACTION field (see 5.2.3.3 and 5.2.1).

### 5.2.3.5 Non-Data NCQ encapsulation

If a command that uses the ZAC Management Out, Non-Data command template is processed in an NCQ environment as subcommand 07h of NCQ NON-DATA command (see ACS-4), then the command inputs are encapsulated as shown in table 18.

**Table 18 — NCQ NON-DATA command encapsulation for the inputs used by the ZAC Management Out command**

NCQ NON-DATA field (see ACS-4)	ZAC Management Out field (see table 17), if any
FEATURE (7:4)	Reserved
FEATURE (15:8)	COUNT (7:0)
COUNT (15:8)	COUNT (15:8)
LBA	LBA
AUXILIARY (31:16)	Reserved
AUXILIARY (15:0)	FEATURE (15:0)

### 5.2.3.6 NCQ encapsulation for a DMA ZAC Management Out command

If a command that uses the ZAC Management Out, DMA command template is processed in an NCQ environment as subcommand 03h of SEND FPDMA QUEUED command (see ACS-4), then the command inputs are encapsulated as shown in table 19.

**Table 19 — SEND FPDMA QUEUED command encapsulation for the inputs used by the ZAC Management Out command**

SEND FPDMA QUEUED field (see ACS-4)	ZAC Management Out field (see table 17), if any
LBA	LBA
AUXILIARY (31:16)	Reserved
AUXILIARY (15:0)	FEATURE (15:0)

## 5.2.4 CLOSE ZONE EXT command – 9Fh/01h, Non-Data

### 5.2.4.1 Feature Set

This 48-bit command is for devices that support the Host Aware Zones feature set (see 4.3) or the Host Managed Zones feature set (see 4.4).

### 5.2.4.2 Description

The CLOSE ZONE EXT command requests that one or more Close Zone functions (see 4.6.3.4.11) be processed.

The CLOSE ZONE EXT command is based on the Non-Data ZAC Management Out command template (see 5.2.3) and is able to use the template's NCQ encapsulation capability.

### 5.2.4.3 Inputs

#### 5.2.4.3.1 Overview

See table 20 for the CLOSE ZONE EXT command inputs.

**Table 20 — CLOSE ZONE EXT command inputs**

Field	Description
FEATURE	<p><b>Bit Description</b></p> <p>15:9 Reserved</p> <p>8 CLOSE ALL bit – See 5.2.4.3.2</p> <p>7:5 Reserved</p> <p>4:0 ZM_ACTION field, shall be set to 01h – See 5.2.2.3.3</p>
COUNT	Reserved
LBA	ZONE ID field – See 5.2.4.3.3
DEVICE	<p><b>Bit Description</b></p> <p>7 Obsolete</p> <p>6 Shall be set to one</p> <p>5 Obsolete</p> <p>4 Transport Dependent – See ACS-4</p> <p>3:0 Reserved</p>
COMMAND	7:0 9Fh

#### 5.2.4.3.2 CLOSE ALL bit

If the CLOSE ALL bit is set to one, the device shall ignore the contents of the ZONE ID field and process a Close Zone function (see 4.6.3.4.11) for each zone with a Zone Condition of IMPLICITLY OPENED or EXPLICITLY OPENED.

If the CLOSE ALL bit is cleared to zero and the zone specified by the ZONE ID field (see 5.2.4.3.3) is in Zone Condition:

- a) IMPLICITLY OPENED, or EXPLICITLY OPENED, then the device shall process a Close Zone function (see 4.6.3.4.11) for the zone specified by the ZONE ID field;
- b) EMPTY, CLOSED, or FULL, then the device shall:
  - A) not change the zone's state; and
  - B) return successful command completion;

- c) READ ONLY, then the device shall:
  - A) return command completion with an error; and
  - B) set the sense key to DATA PROTECT and the additional sense code to ZONE IS READ ONLY (see ACS-4 and SPC-5);or
- d) OFFLINE, then the device shall:
  - A) return command completion with an error; and
  - B) set the sense key to DATA PROTECT and the additional sense code to ZONE IS OFFLINE (see ACS-4 and SPC-5).

#### 5.2.4.3.3 ZONE ID field

The ZONE ID field specifies the lowest LBA of the write pointer zone for which the device shall process the command specified by the ZM\_ACTION field (see 5.2.2.3.3) and the COMMAND field.

If the ZONE ID field does not specify the lowest LBA of the write pointer zone, the device shall:

- a) return command completion with an error; and
- b) set the sense key to ILLEGAL REQUEST and the additional sense code to INVALID FIELD IN CDB (see ACS-4 and SPC-5).

#### 5.2.4.4 Normal Outputs

See table 33.

#### 5.2.4.5 Error Outputs

The device returns command completion with an error if the CLOSE ALL bit (see table 20) is cleared to zero and:

- a) the ZONE ID field does not contain the lowest LBA of a write pointer zone as described in 5.2.4.3.3; or
- b) the Zone Condition is READ ONLY or OFFLINE for the specified zone as described in 5.2.4.3.2.

See table 34.

## 5.2.5 FINISH ZONE EXT command – 9Fh/02h, Non-Data

### 5.2.5.1 Feature Set

This 48-bit command is for devices that support the Host Aware Zones feature set (see 4.3) or the Host Managed Zones feature set (see 4.4).

### 5.2.5.2 Description

The FINISH ZONE EXT command requests that one or more Finish Zone function (see 4.6.3.4.12) be processed.

The device may write the substitute data pattern to the media for unwritten LBAs.

The FINISH ZONE EXT command is based on the Non-Data ZAC Management Out command template (see 5.2.3) and is able to use the template's NCQ encapsulation capability.

### 5.2.5.3 Inputs

#### 5.2.5.3.1 Overview

See table 21 for the FINISH ZONE EXT command inputs.

**Table 21 — FINISH ZONE EXT command inputs**

Field	Description
FEATURE	<p><b>Bit Description</b></p> <p>15:9 Reserved</p> <p>8 FINISH ALL bit – See 5.2.5.3.2</p> <p>7:5 Reserved</p> <p>4:0 ZM_ACTION field, shall be set to 02h – See 5.2.2.3.3</p>
COUNT	Reserved
LBA	ZONE ID field – See 5.2.4.3.3
DEVICE	<p><b>Bit Description</b></p> <p>7 Obsolete</p> <p>6 Shall be set to one</p> <p>5 Obsolete</p> <p>4 Transport Dependent – See ACS-4</p> <p>3:0 Reserved</p>
COMMAND	7:0 9Fh

#### 5.2.5.3.2 FINISH ALL bit

If the FINISH ALL bit is set to one, the device shall ignore the contents of the ZONE ID field and process a Finish Zone function (see 4.6.3.4.12) for each zone with a Zone Condition of IMPLICITLY OPENED, EXPLICITLY OPENED, or CLOSED.

If the FINISH ALL bit is cleared to zero and the zone specified by the ZONE ID field (see 5.2.4.3.3) is in Zone Condition:

- a) EMPTY, IMPLICITLY OPENED, EXPLICITLY OPENED, or CLOSED, then the device shall process a Finish Zone function (see 4.6.3.4.12) for the zone specified by the ZONE ID field;
- b) FULL, then the device shall:



- A) not change the zone's state; and
- B) return successful command completion;
- c) READ ONLY, then the device shall:
  - A) return command completion with an error; and
  - B) set the sense key to DATA PROTECT and the additional sense code to ZONE IS READ ONLY (see ACS-4 and SPC-5);or
- d) OFFLINE, then the device shall:
  - A) return command completion with an error; and
  - B) set the sense key to DATA PROTECT and the additional sense code to ZONE IS OFFLINE (see ACS-4 and SPC-5).

#### 5.2.5.4 Normal Outputs

See table 33.

#### 5.2.5.5 Error Outputs

The device returns command completion with an error if the FINISH ALL bit (see table 21) is cleared to zero and:

- a) the ZONE ID field does not contain the lowest LBA of a write pointer zone as described in 5.2.4.3.3; or
- b) the Zone Condition is READ ONLY or OFFLINE for the specified zone as described in 5.2.5.3.2.

If the processing of any Finish Zone function (see 4.6.3.4.12) returns a failure, the device shall:

- a) return command completion with an error; and
- b) set the sense key to DATA PROTECT and the additional sense code to INSUFFICIENT ZONE RESOURCES (see ACS-4 and SPC-5).

See table 34.

## 5.2.6 OPEN ZONE EXT command – 9Fh/03h, Non-Data

### 5.2.6.1 Feature Set

This 48-bit command is for devices that support the Host Aware Zones feature set (see 4.3) or the Host Managed Zones feature set (see 4.4).

### 5.2.6.2 Description

The OPEN ZONE EXT command requests that one or more Explicitly Open Zone functions (see 4.6.3.4.10) be processed.

The OPEN ZONE EXT command is based on the Non-Data ZAC Management Out command template (see 5.2.3) and is able to use the template's NCQ encapsulation capability.

### 5.2.6.3 Inputs

#### 5.2.6.3.1 Overview

See table 22 for the OPEN ZONE EXT command inputs.

**Table 22 — OPEN ZONE EXT command inputs**

Field	Description
FEATURE	<p><b>Bit Description</b></p> <p>15:9 Reserved</p> <p>8 OPEN ALL bit – See 5.2.6.3.2</p> <p>7:5 Reserved</p> <p>4:0 ZM_ACTION field, shall be set to 03h – See 5.2.2.3.3</p>
COUNT	Reserved
LBA	ZONE ID field – See 5.2.4.3.3
DEVICE	<p><b>Bit Description</b></p> <p>7 Obsolete</p> <p>6 Shall be set to one</p> <p>5 Obsolete</p> <p>4 Transport Dependent – See ACS-4</p> <p>3:0 Reserved</p>
COMMAND	7:0 9Fh

#### 5.2.6.3.2 OPEN ALL bit

If the OPEN ALL bit is set to one, the device shall ignore the contents of the ZONE ID field and process an Explicitly Open Zone function (see 4.6.3.4.10) for each zone with a Zone Condition of CLOSED.

If:

- the device supports the Host Managed Zones feature set (see 4.4);
- the OPEN ALL bit is set to one; and
- the number of zones with a Zone Condition of EXPLICITLY OPENED plus the number of zones with a Zone Condition of CLOSED is greater than the value in the MAXIMUM NUMBER OF OPEN SEQUENTIAL WRITE REQUIRED ZONES field (see 6.2.2.5),

then the device:

- a) shall not process any Explicitly Open Zone functions;
- b) returns command completion with an error (see 5.2.6.5); and
- c) sets the sense key to DATA PROTECT and the additional sense code to INSUFFICIENT ZONE RESOURCES (see 5.2.6.5).

If the OPEN ALL bit is cleared to zero and the zone specified by the ZONE ID field (see 5.2.4.3.3) is in Zone Condition:

- a) EMPTY, IMPLICITLY OPENED, or CLOSED, then the device shall process an Explicitly Open Zone function (see 4.6.3.4.10) for the zone specified by the ZONE ID field;
- b) EXPLICITLY OPENED or FULL, then the device shall:
  - A) not change the zone's state; and
  - B) return successful command completion;
- c) READ ONLY, then the device shall:
  - A) return command completion with an error; and
  - B) set the sense key to DATA PROTECT and the additional sense code to ZONE IS READ ONLY (see ACS-4 and SPC-5);
 or
- d) OFFLINE, then the device shall:
  - A) return command completion with an error; and
  - B) set the sense key to DATA PROTECT and the additional sense code to ZONE IS OFFLINE (see ACS-4 and SPC-5).

#### 5.2.6.4 Normal Outputs

See table 33.

#### 5.2.6.5 Error Outputs

The device returns command completion with an error if the OPEN ALL bit is cleared to zero and:

- a) the ZONE ID field does not contain the lowest LBA of a write pointer zone as described in 5.2.4.3.3; or
- b) the Zone Condition is READ ONLY or OFFLINE for the specified zone as described in 5.2.6.3.2.

If the processing of any Explicitly Open Zone function (see 4.6.3.4.10) returns a failure or too many zones are required to be opened as a result of the OPEN ALL bit being set to one (see 5.2.6.3.2), then the device shall:

- a) return command completion with an error; and
- b) set the sense key to DATA PROTECT and the additional sense code to INSUFFICIENT ZONE RESOURCES (see ACS-4 and SPC-5).

See table 34.

## 5.2.7 REPORT ZONES EXT command – 4Ah/00h, DMA

### 5.2.7.1 Feature Set

This 48-bit command is for devices that support the Host Aware Zones feature set (see 4.3) or the Host Managed Zones feature set (see 4.4).

### 5.2.7.2 Description

The REPORT ZONES EXT command returns the zone structure of a zoned device.

The REPORT ZONES EXT command is based on the ZAC Management In command template (see 5.2.2) and is able to use the template's NCQ encapsulation capability.

### 5.2.7.3 Inputs

#### 5.2.7.3.1 Overview

See table 23 for the REPORT ZONES EXT command inputs.

**Table 23 — REPORT ZONES EXT command inputs**

Field	Description
FEATURE	<p><b>Bit Description</b></p> <ul style="list-style-type: none"> <li>15 PARTIAL bit – See 5.2.7.3.2 and 5.2.7.6.2</li> <li>14 Reserved</li> <li>13:8 REPORTING OPTIONS field – See 5.2.7.3.3</li> <li>7:5 Reserved</li> <li>4:0 ZM_ACTION field, shall be set to 00h – See 5.2.2.3.3</li> </ul>
COUNT	RETURN PAGE COUNT field – See 5.2.2.3.4
LBA	ZONE LOCATOR field – See 5.2.7.3.4
DEVICE	<p><b>Bit Description</b></p> <ul style="list-style-type: none"> <li>7 Obsolete</li> <li>6 Shall be set to one</li> <li>5 Obsolete</li> <li>4 Transport Dependent – See ACS-4</li> <li>3:0 Reserved</li> </ul>
COMMAND	7:0 4Ah

#### 5.2.7.3.2 PARTIAL bit

The PARTIAL bit modifies the definition of the ZONE LIST LENGTH field as described in 5.2.7.6.2.

**5.2.7.3.3 REPORTING OPTIONS field**

The REPORTING OPTIONS field (see table 24) specifies the information to be returned by the REPORT ZONES EXT command.

**Table 24 — REPORTING OPTIONS field**

Code	Description
00h	List all zones in the zoned device
01h	List the zones with a Zone Condition of EMPTY
02h	List the zones with a Zone Condition of IMPLICITLY OPENED
03h	List the zones with a Zone Condition of EXPLICITLY OPENED
04h	List the zones with a Zone Condition of CLOSED
05h	List the zones with a Zone Condition of FULL
06h	List the zones with a Zone Condition of READ ONLY
07h	List the zones with a Zone Condition of OFFLINE
08h..0Fh	Reserved
10h	List the zones with RWP Recommended set to true
11h	List the zones with Non-Sequential Write Resources Active set to true
12h..3Eh	Reserved
3Fh	List the zones with a Zone Condition of NOT WRITE POINTER

**5.2.7.3.4 ZONE LOCATOR field**

The ZONE LOCATOR field should specify the lowest LBA of the first zone to be reported.

If the ZONE LOCATOR field does not specify the lowest LBA of a zone, the first zone to be reported is the zone that contains the LBA specified by the ZONE LOCATOR field.

If the ZONE LOCATOR field specifies an LBA that is greater than the native max address, then:

- a) the command shall be terminated with the ERROR bit set to one and the ID NOT FOUND bit set to one (see ACS-4); and
- b) the sense key shall be set to ILLEGAL REQUEST and the additional sense code shall be set to LOGICAL BLOCK ADDRESS OUT OF RANGE (see ACS-4 and SPC-5).

**5.2.7.4 Normal Outputs**

See table 33.

**5.2.7.5 Error Outputs**

See table 34.

## 5.2.7.6 Input From the Device to the Host Data Structure

### 5.2.7.6.1 Overview

The format of the data input from the device to the host (see table 25) includes a zone descriptors list that contains zone descriptor 0 through zone descriptor N.

**Table 25 — REPORT ZONES EXT input from device to host**

Offset	Type	Description
0..3	DWord	ZONE LIST LENGTH field (see 5.2.7.6.2)
4	Byte	<b>Bits Description</b> 7:4 Reserved 3:0 SAME field (see 5.2.7.6.3)
5..7		Reserved
8..15	QWord	<b>Bits Description</b> 63:48 Reserved 47:0 MAXIMUM LBA field (see 5.2.7.6.4)
16..63		Reserved
64..127	Bytes	Zone descriptor 0 (see 5.2.7.6.6)
128..191	Bytes	Zone descriptor 1
...		...
s-63..s	Bytes	Last zone descriptor N
s+1..(512 × pc-1)	Bytes	Zero padding (see 5.2.2.3.4)

### 5.2.7.6.2 ZONE LIST LENGTH field

The ZONE LIST LENGTH field contains the length in bytes of the zone descriptors list (i.e., bytes 63..s in table 25). Based on the value of the PARTIAL bit (see 5.2.7.3.2), the zone descriptors in the zone descriptors list are all of the zone descriptors or a subset of all the zone descriptors associated with each zone specified by the REPORTING OPTIONS field (see 5.2.7.3.3) that includes an LBA that is greater than or equal to the LBA specified by the ZONE LOCATOR field (see 5.2.7.3.4).

If the PARTIAL bit (see table 23) is:

- a) cleared to zero, then the ZONE LIST LENGTH field shall:
  - A) contain the length in bytes of the zone descriptors list (i.e., s-63); and
  - B) not be altered based on the RETURN PAGE COUNT field (see 5.2.2.3.4);
 or
- b) set to one, then the ZONE LIST LENGTH field shall contain the lesser of:
  - A) 64 multiplied by the number of zones that:
    - a) include a lowest LBA that is greater than or equal to the lowest LBA of the zone specified by the ZONE LOCATOR field (see 5.2.7.3.4); and
    - b) match the criteria specified by the REPORTING OPTIONS field (see 5.2.7.3.3);
 or
  - B) (contents of the RETURN PAGE COUNT field × 512) – 64.

The zone list length is the number of zones reported multiplied by 64.

**5.2.7.6.3 SAME field**

The SAME field (see table 26) indicates same zone type and same zone length relationships between the zone described by zone descriptor 0 (see table 25) and the zones described by the other zone descriptors in the zone descriptors list in the input from the device to the host. If the ZONE LIST LENGTH field is zero then the SAME field is invalid and should be ignored by the host.

**Table 26 — SAME field**

<b>Code</b>	<b>Description</b>
0h	The zone type and zone length in each zone descriptor in the zone descriptors list may be different.
1h	The zone type and zone length in each zone descriptor in the zone descriptors list are equal to the zone type and zone length indicated by zone descriptor 0.
2h	The zone type in each zone descriptor in the zone descriptors list is equal to the zone type indicated by zone descriptor 0. The zone length of each zone except the last zone is equal to the zone length indicated by zone descriptor 0. The zone length of the last zone is different than the zone length of the other zones.
3h	The zone length in each zone descriptor in the zone descriptors list is equal to the zone length indicated by zone descriptor for the first zone on the zoned device (i.e., the zone that starts at LBA 0). The zone types for each descriptor may be different.
all others	Reserved

**5.2.7.6.4 MAXIMUM LBA field**

The MAXIMUM LBA field indicates the LBA of the last logical sector on the device, including all logical sectors in all zones.

**5.2.7.6.5 Zero padding**

The total bytes input from the device to the host shall be:

$$512 \times pc$$

where:

pc is the contents of the RETURN PAGE COUNT field (see 5.2.2.3.4) in the command inputs (see table 23).

The device shall insert as many bytes cleared to zero after the last zone descriptor as are required to make total bytes input from the device to the host equal of  $512 \times pc$ .

### 5.2.7.6.6 Zone descriptor format

#### 5.2.7.6.6.1 Overview

The zone descriptors list contains zone descriptors that shall be sorted in ascending order based on the lowest LBA of each zone. Each zone descriptor (see table 27) describes one zone. A nonzero value in the SAME field (see 5.2.7.6.3) indicates which fields in which descriptors contain a value that is the same as the equivalent value in zone descriptor 0.

If the PARTIAL bit is cleared to zero (see 5.2.7.3.2), there is one zone descriptor for each zone with:

- a) a lowest LBA that is greater than or equal to the lowest LBA of the zone specified by the ZONE LOCATOR field (see 5.2.7.3.4); and
- b) attributes that match those specified by the REPORTING OPTIONS field (see 5.2.7.3.3).

**Table 27 — Zone descriptor format**

Offset	Type	Description
0	Byte	<b>Bit Meaning</b> 7:4 Reserved 3:0 ZONE TYPE field (see 5.2.7.6.6.2)
1	Byte	<b>Bit Meaning</b> 7:4 ZONE CONDITION field (see 5.2.7.6.6.3) 3:2 Reserved 1 NON_SEQ bit (see 5.2.7.6.6.4) 0 RESET bit (see 5.2.7.6.6.5)
2..7	Bytes	Reserved
8..15	QWord	<b>Bits Description</b> 63:48 Reserved 47:0 ZONE LENGTH field (see 5.2.7.6.6.6)
16..23	QWord	<b>Bits Description</b> 63:48 Reserved 47:0 ZONE START LBA field (see 5.2.7.6.6.7)
24..31	QWord	<b>Bits Description</b> 63:48 Reserved 47:0 WRITE POINTER LBA field (see 5.2.7.6.6.8)
32..63	Bytes	Reserved

#### 5.2.7.6.6.2 ZONE TYPE field

The ZONE TYPE field (see table 28) indicates the Zone Type (see 4.5.2).

**Table 28 — ZONE TYPE field**

Code	Zone Type
0h	Reserved
1h	CONVENTIONAL
2h	SEQUENTIAL WRITE REQUIRED
3h	SEQUENTIAL WRITE PREFERRED
4h..Fh	Reserved



**5.2.7.6.6.3 ZONE CONDITION field**

The ZONE CONDITION field (see table 29) indicates the Zone Condition (see 4.5.3).

**Table 29 — ZONE CONDITION field**

Code	Zone Condition
0h	NOT WRITE POINTER
1h	EMPTY
2h	IMPLICITLY OPENED
3h	EXPLICITLY OPENED
4h	CLOSED
5h .. Ch	Reserved
Dh	READ ONLY
Eh	FULL
Fh	OFFLINE

**5.2.7.6.6.4 NON\_SEQ bit**

The value of the non-sequential (NON\_SEQ) bit is based on the Non-Sequential Write Resources Active zone attribute (see 4.5.6). If Non-Sequential Write Resources Active is:

- a) false, then the NON\_SEQ bit shall be cleared to zero; or
- b) true, then the NON\_SEQ bit shall be set to one.

**5.2.7.6.6.5 RESET bit**

The value of the RESET bit is based on the RWP Recommended zone attribute (see 4.5.5). If RWP Recommended is:

- a) false, then the RESET bit shall be cleared to zero; or
- b) true, then the RESET bit shall be set to one.

**5.2.7.6.6.6 ZONE LENGTH field**

The ZONE LENGTH field indicates the number of logical sectors in the zone described by this zone descriptor.

**5.2.7.6.6.7 ZONE START LBA field**

The ZONE START LBA field indicates the lowest LBA in the zone described by this zone descriptor.

**5.2.7.6.6.8 WRITE POINTER LBA field**

The WRITE POINTER LBA field indicates the starting LBA that the host should specify in the next write command associated with the zone described by this zone descriptor.

The host should ignore the contents of the WRITE POINTER LBA field if the Zone Condition indicates that the write pointer is invalid (see table 4 and table 7).

## 5.2.8 RESET WRITE POINTER EXT command – 9Fh/04h, Non-Data

### 5.2.8.1 Feature Set

This 48-bit command is for devices that support the Host Aware Zones feature set (see 4.3) or the Host Managed Zones feature set (see 4.4).

### 5.2.8.2 Description

The RESET WRITE POINTER EXT command requests that one or more Reset Write Pointer functions (see 4.6.3.4.13) be processed.

The RESET WRITE POINTER EXT command is based on the Non-Data ZAC Management Out command template (see 5.2.3) and is able to use the template's NCQ encapsulation capability.

### 5.2.8.3 Inputs

#### 5.2.8.3.1 Overview

See table 30 for the RESET WRITE POINTER EXT command inputs.

**Table 30 — RESET WRITE POINTER EXT command inputs**

Field	Description
FEATURE	<p><b>Bit Description</b></p> <p>15:9 Reserved</p> <p>8 RESET ALL bit – See 5.2.8.3.2</p> <p>7:5 Reserved</p> <p>4:0 ZM_ACTION field, shall be set to 04h – See 5.2.2.3.3</p>
COUNT	Reserved
LBA	ZONE ID field – See 5.2.4.3.3
DEVICE	<p><b>Bit Description</b></p> <p>7 Obsolete</p> <p>6 Shall be set to one</p> <p>5 Obsolete</p> <p>4 Transport Dependent – See ACS-4</p> <p>3:0 Reserved</p>
COMMAND	7:0 9Fh

#### 5.2.8.3.2 RESET ALL bit

If the RESET ALL bit is set to one, the device shall ignore the contents of the ZONE ID field and process a Reset Write Pointer function (see 4.6.3.4.13) for each zone with a Zone Condition of IMPLICITLY OPENED, EXPLICITLY OPENED, CLOSED, or FULL.

If the RESET ALL bit is cleared to zero and the zone specified by the ZONE ID field (see 5.2.4.3.3) is in Zone Condition:

- a) IMPLICITLY OPENED, EXPLICITLY OPENED, CLOSED, or FULL, then the device shall process a Reset Write Pointer function (see 4.6.3.4.13) for the zone specified by the ZONE ID field;
- b) EMPTY, then the device shall:
  - A) not change the zone's state; and
  - B) return successful command completion;

- c) READ ONLY, then the device shall:
  - A) return command completion with an error; and
  - B) set the sense key to DATA PROTECT and the additional sense code to ZONE IS READ ONLY (see ACS-4 and SPC-5);or
- d) OFFLINE, then the device shall:
  - A) return command completion with an error; and
  - B) set the sense key to DATA PROTECT and the additional sense code to ZONE IS OFFLINE (see ACS-4 and SPC-5).

#### 5.2.8.4 Normal Outputs

See table 33.

#### 5.2.8.5 Error Outputs

The device returns command completion with an error if the RESET ALL bit is cleared to zero and:

- a) the ZONE ID field does not contain the lowest LBA of a write pointer zone as described in 5.2.4.3.3; or
- b) the Zone Condition is READ ONLY, or OFFLINE for the specified zone as described in 5.2.8.3.2.

See table 34.

## 6 Log definitions

### 6.1 Introduction

Clause 6 describes the logs defined by this standard in addition to the logs defined by ACS-4. The logs defined by this standard are part of the Host Aware Zones feature set (see 4.3) and the Host Managed Zones feature set (see 4.4).

Unless otherwise defined, all aspects of the log descriptions in clause 6 are as defined in ACS-4.

### 6.2 IDENTIFY DEVICE data log (Log Address 30h)

#### 6.2.1 Overview

With the exception of the Zoned Device Information page (see 6.2.2), everything related to the IDENTIFY DEVICE data log is defined in ACS-4.

#### 6.2.2 Zoned Device Information (page 09h)

##### 6.2.2.1 Overview

The Zoned Device Information log page (see table 31) provides information about the capabilities and functional settings of a zoned block device.

**Table 31 — Zoned Device Information (part 1 of 2)**

Offset	Type	Content
0..7	QWord	Zoned Device Information page information header.
		<p><b>Bit Meaning</b></p> <p>63 Shall be set to one.</p> <p>62:24 Reserved</p> <p>23:16 Page Number. Shall be set to 09h.</p> <p>15:0 Revision number. Shall be set to 0001h.</p>
8..15	QWord	Zoned device capabilities
		<p><b>Bit Meaning</b></p> <p>63 Shall be set to one.</p> <p>62:1 Reserved</p> <p>0 URSWRZ bit (see 6.2.2.2)</p>
16..23	QWord	Zoned device settings
		<p><b>Bit Meaning</b></p> <p>63 Shall be set to one.</p> <p>62:0 Reserved</p>
24..31	QWord	Optimal number of open Sequential Write Preferred zones
		<p><b>Bit Meaning</b></p> <p>63 Shall be set to one.</p> <p>62:32 Reserved</p> <p>31:0 OPTIMAL NUMBER OF OPEN SEQUENTIAL WRITE PREFERRED ZONES field (see 6.2.2.3)</p>

Table 31 — Zoned Device Information (part 2 of 2)

Offset	Type	Content
32..39	QWord	Optimal number of non-sequentially written Sequential Write Preferred zones
		<p><b>Bit Meaning</b></p> <p>63 Shall be set to one.</p> <p>62:32 Reserved</p> <p>31:0 OPTIMAL NUMBER OF NON-SEQUENTIALLY WRITTEN SEQUENTIAL WRITE PREFERRED ZONES field (see 6.2.2.4)</p>
40..47	QWord	Maximum number of open Sequential Write Required zones
		<p><b>Bit Meaning</b></p> <p>63 Shall be set to one.</p> <p>62:32 Reserved</p> <p>31:0 MAXIMUM NUMBER OF OPEN SEQUENTIAL WRITE REQUIRED ZONES field (see 6.2.2.5)</p>
48..55	QWord	Version information
		<p><b>Bit Meaning</b></p> <p>63 Shall be set to one.</p> <p>62:16 Reserved</p> <p>15:0 ZAC MINOR VERSION field (see 6.2.2.6.1)</p>
56..511		Reserved

### 6.2.2.2 URSWRZ bit

An unrestricted read in sequential write required zone (URSWRZ) bit cleared to zero indicates that the device does not support reading unwritten logical sectors or reads across zone boundaries in Sequential Write Required zones. A URSWRZ bit set to one indicates that the device supports reading unwritten logical sectors and reads across zone boundaries in Sequential Write Required zones.

The details of processing reads in Sequential Write Required zones are described in 4.6.3.3.2.

### 6.2.2.3 OPTIMAL NUMBER OF OPEN SEQUENTIAL WRITE PREFERRED ZONES field

The OPTIMAL NUMBER OF OPEN SEQUENTIAL WRITE PREFERRED ZONES field indicates the recommended upper limit for the number of Sequential Write Preferred zones that may have a Zone Condition of IMPLICITLY OPENED or EXPLICITLY OPENED. The Manage Open Zone Resources function may enforce this recommendation as described in 4.6.3.4.9.

A value of FFFF FFFFh indicates that the optimal number of open Sequential Write Preferred zones is not reported.

### 6.2.2.4 OPTIMAL NUMBER OF NON-SEQUENTIALLY WRITTEN SEQUENTIAL WRITE PREFERRED ZONES field

The OPTIMAL NUMBER OF NON-SEQUENTIALLY WRITTEN SEQUENTIAL WRITE PREFERRED ZONES field indicates the recommended upper limit for the number of Sequential Write Preferred zones in which at least one non-sequential write commands have been processed (i.e., the recommended upper limit for the number of zones with Non-Sequential Write Resources Active true).

A value of FFFF FFFFh indicates that the optimal number of non-sequentially written Sequential Write Preferred zones is not reported.

### 6.2.2.5 MAXIMUM NUMBER OF OPEN SEQUENTIAL WRITE REQUIRED ZONES field

The MAXIMUM NUMBER OF OPEN SEQUENTIAL WRITE REQUIRED ZONES field indicates the maximum number of Sequential Write Required zones that are allowed to have a Zone Condition of EXPLICITLY OPENED. The Manage Open Zone Resources function enforces this limit as described in 4.6.3.4.9.

A value of FFFF FFFFh indicates that there is no limit for the number of open Sequential Write Required zones.

### 6.2.2.6 Version information

#### 6.2.2.6.1 ZAC MINOR VERSION field

The ZAC MINOR VERSION field (see table 32) indicates the version of the standard that guided the device implementation.

**Table 32 — ZAC MINOR VERSION field**

<b>Value</b>	<b>Minor version</b>
0000h	ZAC minor version is not reported
0000h..05CEh	Reserved
05CFh	ZAC revision 05
05D0h..A36Bh	Reserved
A36Ch	ZAC revision 04
A36Dh..B6E7h	Reserved
B6E8h	ZAC revision 01
B6E9h..FFFEh	Reserved
FFFFh	ZAC minor version is not reported

## 7 Normal and Error Outputs

### 7.1 Overview

The commands listed in clause 5 each have subclauses labeled Normal Outputs and Error Outputs. Subclauses 7.2 and 7.3 document the return data format for all the commands described in clause 5. Each command in clause 5 may provide additional information about a normal or error output, however, all the information specified in clause 7 shall also apply to the command.

The references preceding each table indicate each command that generates the output in the table.

### 7.2 Normal Outputs

The tables in this subclause specify the Normal Outputs a command returns.

Table 33 specifies the normal outputs for the commands defined in 5.2.4, 5.2.5, 5.2.6, 5.2.7, and 5.2.8.

**Table 33 — Generic ZAC Extended Normal Output**

Field	Description
ERROR	Reserved
COUNT	Reserved
LBA	Reserved
DEVICE	<p><b>Bit Description</b></p> <ul style="list-style-type: none"> <li>7 Obsolete</li> <li>6 N/A</li> <li>5 Obsolete</li> <li>4 Transport Dependent – See ACS-4</li> <li>3:0 Reserved</li> </ul>
STATUS	<p><b>Bit Description</b></p> <ul style="list-style-type: none"> <li>7:6 Transport Dependent – See ACS-4</li> <li>5 DEVICE FAULT bit – See ACS-4</li> <li>4 N/A</li> <li>3 Transport Dependent – See ACS-4</li> <li>2 N/A or ALIGNMENT ERROR bit – See ACS-4</li> <li>1 SENSE DATA AVAILABLE bit – See ACS-4</li> <li>0 ERROR bit – See ACS-4</li> </ul>

### 7.3 Error Outputs

The tables in this subclause specify the Error Outputs a command returns. References to these tables are found in clause 5.

If the Sense Data Reporting feature set is enabled and there is sense data available, then the ERROR field shall be set to 7Fh and the ERROR bit in the STATUS field shall be set to one. If the Sense Data Reporting feature set has been enabled with the SENSE DATA ENABLED bit set to one (see ACS-4), then the device notifies the host of additional information by setting the SENSE DATA AVAILABLE bit in the STATUS field to one.

Table 34 specifies the error outputs for the commands defined in 5.2.4, 5.2.5, 5.2.6, 5.2.7, and 5.2.8.

**Table 34 — Generic ZAC Extended Error Output**

Field	Description
ERROR	<p><b>Bit Description</b></p> <p>7:3 N/A</p> <p>2 ABORT bit – See ACS-4</p> <p>1:0 N/A</p>
COUNT	Reserved
LBA	Reserved
DEVICE	<p><b>Bit Description</b></p> <p>7 Obsolete</p> <p>6 N/A</p> <p>5 Obsolete</p> <p>4 Transport Dependent – See ACS-4</p> <p>3:0 Reserved</p>
STATUS	<p><b>Bit Description</b></p> <p>7:6 Transport Dependent – See ACS-4</p> <p>5 DEVICE FAULT bit – See ACS-4</p> <p>4 N/A</p> <p>3 Transport Dependent – See ACS-4</p> <p>2 N/A</p> <p>1 SENSE DATA AVAILABLE bit – See ACS-4</p> <p>0 ERROR bit – See ACS-4</p>



## Annex A (informative)

### Host Considerations for Zoned Devices

#### A.1 Introduction

Zoned devices add host management requirements to the ACS-4 requirements associated with devices. One perspective on the added zoned device requirements is found in the article *Shingled Magnetic Recording: Areal Density Increase Requires New Data Management* (see Bibliography) that discusses:

- a) Cooperatively Managed SMR concepts that are related to the Host Aware Zones feature set (see 4.3);
- b) Host-Managed SMR concepts that are related to the Host Managed Zones feature set (see 4.4); and
- c) Drive-Managed SMR concepts that are outside the scope of this standard.

To facilitate better interactions between hosts and devices, this annex considers descriptions of ZAC features from a host perspective.

#### A.2 Writing to write pointer zones

In a write pointer zone (see 4.6.3), the preferred place to write is at the valid write pointer. In a Sequential Write Required zone (see 4.6.3.2), writes are required to begin at the valid write pointer.

Although a zone with a Zone Condition of FULL has an invalid write pointer, the device is able to process write commands in Sequential Write Preferred zones (see 4.6.3.2) with a Zone Condition of FULL.

While a zone is an open zone, the write pointer may be volatile.

#### A.3 Open zone considerations

##### A.3.1 Overview

A zoned device is able to process write commands only on LBAs in:

- a) open zones; or
- b) Sequential Write Preferred zones (see 4.6.3.2) with a Zone Condition of FULL.

Open zones have access to open zone resources to support writing that are not available to zones that are not open. Limitations on the availability of open zone resources and the number of zones that are allowed to be open are indicated by:

- a) the OPTIMAL NUMBER OF OPEN SEQUENTIAL WRITE PREFERRED ZONES field (see A.4.1); and
- b) the MAXIMUM NUMBER OF OPEN SEQUENTIAL WRITE REQUIRED ZONES field (see A.4.2).

Zones with a Zone Condition of EMPTY or CLOSED are opened as a result of processing:

- a) a write command (see ACS-4);
- b) a FINISH ZONE EXT command (see 5.2.5);
- c) an OPEN ZONE EXT command (see 5.2.6); or
- d) other identified ACS-4 commands (e.g., SECURITY ERASE UNIT commands, Sanitize commands).

If a write command is processed on LBAs in a zone with a Zone Condition of EMPTY or CLOSED, the Zone Condition state machine (see 4.6.3.4):

- 1) processes a Manage Open Zone Resources function (see 4.6.3.4.9);
- 2) opens the zone, if open zone resources are available; and

- 3) processes the write command.

A write command results in an error if processed on any zone with a Zone Condition of READ ONLY, OFFLINE, or a Sequential Write Required zone with a Zone Condition of FULL. If a write command specifies LBAs in a zone with a Zone Condition of EMPTY or CLOSED, then the device:

- 1) performs a Manage Open Zone Resources function (see 4.6.3.4.9);
- 2) opens the zone, if open zone resources are available; and
- 3) performs the write command.

An open zone remains open until:

- a) a write command is processed for this zone that writes from the write pointer to the highest LBA in the zone;
- b) this zone:
  - A) has a Zone Condition of IMPLICITLY OPENED; and
  - B) a Manage Open Zone Resources function (see 4.6.3.4.9) processed for a different zone selects and closes this zone;
- c) a Finish Zone function (see 4.6.3.4.12 and A.3.4) is processed for this zone;
- d) a Reset Write Pointer function (see 4.6.3.4.13) is processed for this zone;
- e) a Close Zone function (see 4.6.3.4.11) is processed for this zone;
- f) this zone's Zone Condition becomes READ ONLY or OFFLINE; or
- g) a power-on reset is processed.

Until a zone is closed, the following may be volatile:

- a) user data written in that zone; and
- b) the write pointer, if any.

### A.3.2 Explicitly opened zones and implicitly opened zones

An open zone is a zone with a Zone Condition of IMPLICITLY OPENED or EXPLICITLY OPENED. For any zone with a Zone Condition of EMPTY or CLOSED, a device opens a zone with a Zone Condition of:

- a) EXPLICITLY OPENED as a result of processing an Explicitly Open Zone function (see 4.6.3.4.10); or
- b) IMPLICITLY OPENED as a result of processing:
  - A) a write command;
  - B) a Finish Zone function (see 4.6.3.4.12 and A.3.4); or
  - C) other identified ACS-4 commands (e.g., SECURITY ERASE UNIT commands, Sanitize commands).

If a device processes an Explicitly Open Zone function on a zone with a Zone Condition of IMPLICITLY OPENED, the result is a zone with a Zone Condition of EXPLICITLY OPENED.

The Manage Open Zone Resources function (see 4.6.3.4.9) is:

- a) allowed to close a zone with a Zone Condition of IMPLICITLY OPENED; and
- b) not allowed to close a zone with a Zone Condition of EXPLICITLY OPENED.

As a result, a host may send an OPEN ZONE EXT command (see 5.2.6) for a zone in which host timing requirements are unable to tolerate the unplanned overhead, if any, associated with reopening a zone with a Zone Condition of IMPLICITLY OPENED that a Manage Open Zone Resources function selected to be closed.

After a power-on reset, no zones have a Zone Condition of IMPLICITLY OPENED or EXPLICITLY OPENED. For any zone that was an open zone before the power-on reset, after the power-on reset the Zone Condition is:

- a) EMPTY, if the write pointer indicates the lowest LBA in the zone and Non-Sequential Resources Active is false;
- b) CLOSED, if the write pointer does not indicate the lowest LBA in the zone or Non-Sequential Resources Active is true; or
- c) READ ONLY or OFFLINE, if the device detects an applicable condition (e.g., a media failure (see 4.7.4)) during the processing of the power-on reset.

### A.3.3 Opening and closing zones

A host uses one of the following techniques to open a zone for writing and later to close that zone:

- a) an OPEN ZONE EXT command (see 5.2.6) to open a zone and a CLOSE ZONE EXT command (see 5.2.4) to close that zone; or
- b) a write command to open a zone and:
  - A) a CLOSE ZONE EXT command to close that zone; or
  - B) allowing the Manage Open Zone Resources function (see 4.6.3.4.9) to close that zone.

A host may use different techniques to manage the opening and closing of different zones.

A host may use CLOSE ZONE EXT commands to close more zones than the number necessary to satisfy limitations on the number of zones that are allowed to be open (see A.3.1).

### A.3.4 Finish Zone function considerations

A zone may be opened (e.g., with an OPEN ZONE EXT command (see 5.2.6)) prior to processing a Finish Zone function (see 4.6.3.4.12).

Regardless of how a zone is opened before processing a Finish Zone function, completion of that Finish Zone function changes the Zone Condition to FULL, with the result that open zone resources become available to open a different zone.

## A.4 Open zone resources considerations based on zone type

### A.4.1 Sequential Write Preferred zones

A host:

- a) reads the OPTIMAL NUMBER OF OPEN SEQUENTIAL WRITE PREFERRED ZONES field (see 6.2.2.3) to determine the recommended number of zones open at the same time; and
- b) sends:
  - A) a REPORT ZONES EXT command (see 5.2.7) with:
    - a) the REPORTING OPTIONS field (see 5.2.7.3.3) set to 02h; and
    - b) the PARTIAL bit (see 5.2.7.3.2) cleared to zero, to return the list of zones with a Zone Condition of IMPLICITLY OPENED;
 and
  - B) a REPORT ZONES EXT command with:
    - a) the REPORTING OPTIONS field set to 03h; and
    - b) the PARTIAL bit cleared to zero, to return the list of zones with a Zone Condition of EXPLICITLY OPENED, to determine the number of open zones (i.e., zones with a Zone Condition of IMPLICITLY OPENED plus zones with a Zone Condition of EXPLICITLY OPENED).

If the number of open zones is equal to or greater than the optimal number of open Sequential Write Preferred zones, then:

- a) the preferred host behavior is to write to a zone with Zone Condition of EXPLICITLY OPENED or IMPLICITLY OPENED instead of writing to a zone with Zone Condition of EMPTY or CLOSED; or
- b) before writing to a zone with Zone Condition of EMPTY or CLOSED, the host may send a CLOSE ZONE EXT command (see 5.2.4) to specify the zone to be closed instead of allowing the device to select which zone, if any, to close when the write occurs.

A host:

- a) reads the OPTIMAL NUMBER OF NON-SEQUENTIALLY WRITTEN SEQUENTIAL WRITE PREFERRED ZONES field (see 6.2.2.4) to determine the optimum number of zones that are allowed to be non-sequentially written at the same time; and
- b) sends a REPORT ZONES EXT command (see 5.2.7) with:

- A) the REPORTING OPTIONS field (see 5.2.7.3.3) set to 11h; and
  - B) the PARTIAL bit (see 5.2.7.3.2) cleared to zero,
- to determine the list of zones with Non-Sequential Write Resources Active set to true.

The preferred host behavior is to manage writes to a zoned device to ensure that the number of Sequential Write Preferred zones with Non-Sequential Write Resources Active set to true is less than the optimal number of non-sequentially written sequential write preferred zones.

#### A.4.2 Sequential Write Required zones

The value in the MAXIMUM NUMBER OF OPEN SEQUENTIAL WRITE REQUIRED ZONES field (see 6.2.2.5) indicates the allowed maximum number of zones with a Zone Condition of EXPLICITLY OPENED. If this maximum has been reached and all zones with a Zone Condition of IMPLICITLY OPENED have been closed by the Manage Open Zone Resources function (see 4.6.3.4.9), then the device is required to complete any command with an error if that command requires a zone to be opened (e.g., a write command, an OPEN ZONE EXT command (see 5.2.6)).

A host:

- a) reads the MAXIMUM NUMBER OF OPEN SEQUENTIAL WRITE REQUIRED ZONES field (see 6.2.2.5) to determine the recommended number of zones open at the same time; and
- b) sends:
  - A) a REPORT ZONES EXT command (see 5.2.7) with:
    - a) the REPORTING OPTIONS field (see 5.2.7.3.3) set to 02h; and
    - b) the PARTIAL bit (see 5.2.7.3.2) cleared to zero,
 to return the list of zones with a Zone Condition of IMPLICITLY OPENED;
  - and
  - B) a REPORT ZONES EXT command with:
    - a) the REPORTING OPTIONS field set to 03h; and
    - b) the PARTIAL bit cleared to zero,
 to return the list of zones with a Zone Condition of EXPLICITLY OPENED,
 to determine the number of open zones (i.e., zones with a Zone Condition of IMPLICITLY OPENED plus zones with a Zone Condition of EXPLICITLY OPENED).

If the number of open zones is equal to or greater than the maximum number of open Sequential Write Required zones, then:

- a) the preferred host behavior is to write to a zone with Zone Condition of EXPLICITLY OPENED or IMPLICITLY OPENED instead of writing to a zone with Zone Condition of EMPTY or CLOSED; or
- b) to avoid commands that are completed with an error when writing to a zone with Zone Condition of EMPTY or CLOSED, the host sends CLOSE ZONE EXT command (see 5.2.4) to specify a zone to be closed.

## A.5 Partial failures

### A.5.1 Overview

The Zone Conditions READ ONLY and OFFLINE provide the ability for hosts to continue using a device after parts of its storage capacity have stopped operating (e.g., the use of the OFFLINE Zone Condition as a response to a media failures (see 4.7.4)).

Many of the conditions that result in the Zone Condition becoming READ ONLY and OFFLINE are outside the scope of this standard.

### A.5.2 Sanitize considerations

Zoned devices may change the Zone Condition to OFFLINE in any zone that has a Zone Condition of READ ONLY prior to the processing of a sanitize operation (see ACS-4).

Zoned devices may avoid failing a sanitize operation by changing the Zone Condition to OFFLINE in any zone where ACS-4 requirements for media failures (see 4.7.4) prevent successful completion of a sanitize operation.

## Bibliography

USENIX ;login:, v 38, n 3, *Shingled Magnetic Recording: Areal Density Increase Requires New Data Management*, Tim Feldman, Garth Gibson, June 2013<sup>1 2</sup>

---

<sup>1</sup> See <https://www.usenix.org/publications/login>.

<sup>2</sup> Reprints are available from, <http://www.cs.cmu.edu/~./garth/#pubs>.