



BSI Standards Publication

Information technology — Small computer system interface (SCSI)

Part 346: Zoned Block Commands — 2 (ZBC-2)

National foreword

This British Standard is the UK implementation of ISO/IEC 14776-346:2024.

The UK participation in its preparation was entrusted to Technical Committee ICT/1, Information systems co-ordination.

A list of organizations represented on this committee can be obtained on request to its committee manager.

Contractual and legal considerations

This publication has been prepared in good faith, however no representation, warranty, assurance or undertaking (express or implied) is or will be made, and no responsibility or liability is or will be accepted by BSI in relation to the adequacy, accuracy, completeness or reasonableness of this publication. All and any such responsibility and liability is expressly disclaimed to the full extent permitted by the law.

This publication is provided as is, and is to be used at the recipient's own risk.

The recipient is advised to consider seeking professional guidance with respect to its use of this publication.

This publication is not intended to constitute a contract. Users are responsible for its correct application.

© The British Standards Institution 2024
Published by BSI Standards Limited 2024

ISBN 978 0 539 28384 6

ICS 35.200

Compliance with a British Standard cannot confer immunity from legal obligations.

This British Standard was published under the authority of the Standards Policy and Strategy Committee on 30 November 2024.

Amendments/corrigenda issued since publication

Date	Text affected



International
Standard

ISO/IEC 14776-346

**Information technology — Small
computer system interface (SCSI) —**

**Part 346:
Zoned Block Commands - 2 (ZBC-2)**

*Technologies de l'information — Interface de petit système
d'ordinateur (SCSI) —*

Partie 346: Commandes de blocs de zones - 2 (ZBC-2)

**First edition
2024-09**

Reference number
ISO/IEC 14776-346:2024(en)

Full document details: BS ISO/IEC 14776-346:2024

 **Documents**
© ISO/IEC 2024 Document Sharing Hub



COPYRIGHT PROTECTED DOCUMENT

© ISO/IEC 2024

All rights reserved. Unless otherwise specified, or required in the context of its implementation, no part of this publication may be reproduced or utilized otherwise in any form or by any means, electronic or mechanical, including photocopying, or posting on the internet or an intranet, without prior written permission. Permission can be requested from either ISO at the address below or ISO's member body in the country of the requester.

ISO copyright office
CP 401 • Ch. de Blandonnet 8
CH-1214 Vernier, Geneva
Phone: +41 22 749 01 11
Email: copyright@iso.org
Website: www.iso.org

Published in Switzerland

Foreword

ISO (the International Organization for Standardization) and IEC (the International Electrotechnical Commission) form the specialized system for worldwide standardization. National bodies that are members of ISO or IEC participate in the development of International Standards through technical committees established by the respective organization to deal with particular fields of technical activity. ISO and IEC technical committees collaborate in fields of mutual interest. Other international organizations, governmental and non-governmental, in liaison with ISO and IEC, also take part in the work.

The procedures used to develop this document and those intended for its further maintenance are described in the ISO/IEC Directives, Part 1. In particular, the different approval criteria needed for the different types of document should be noted (see www.iso.org/directives or www.iec.ch/members_experts/refdocs).

ISO and IEC draw attention to the possibility that the implementation of this document may involve the use of (a) patent(s). ISO and IEC take no position concerning the evidence, validity or applicability of any claimed patent rights in respect thereof. As of the date of publication of this document, ISO and IEC had not received notice of (a) patent(s) which may be required to implement this document. However, implementers are cautioned that this may not represent the latest information, which may be obtained from the patent database available at www.iso.org/patents and <https://patents.iec.ch>. ISO and IEC shall not be held responsible for identifying any or all such patent rights.

Any trade name used in this document is information given for the convenience of users and does not constitute an endorsement.

For an explanation of the voluntary nature of standards, the meaning of ISO specific terms and expressions related to conformity assessment, as well as information about ISO's adherence to the World Trade Organization (WTO) principles in the Technical Barriers to Trade (TBT), see www.iso.org/iso/foreword.html. In the IEC, see www.iec.ch/understanding-standards.

This document was prepared by INCITS (as INCITS 550-2023) and drafted in accordance with its editorial rules. It was assigned to Joint Technical Committee ISO/IEC JTC 1, *Information technology*, and adopted under the “fast-track procedure”.

A list of all parts in the ISO/IEC 14776 series can be found on the ISO and IEC websites.

Any feedback or questions on this document should be directed to the user's national standards body. A complete listing of these bodies can be found at www.iso.org/members.html and www.iec.ch/national-committees.

Contents

	Page
FOREWORD.....	xii
INTRODUCTION	xiii
General	xiii
SCSI standards family	xiii
1 Scope	1
2 Normative references	1
3 Definitions, symbols, abbreviations, and conventions	2
3.1 Definitions	2
3.2 Symbols and abbreviations	8
3.2.1 Abbreviations	8
3.2.2 Mathematical operators	8
3.3 Keywords	8
3.4 Editorial conventions	10
3.5 Numeric and character conventions	10
3.5.1 Numeric conventions	10
3.5.2 Units of measure	11
3.6 Bit and byte ordering	12
3.7 Notation for state diagrams	14
4 Zoned Block Device Model	15
4.1 Zoned Block Device model overview	15
4.1.1 Established SCSI concepts	15
4.1.2 Peripheral device type and supported commands	16
4.2 Zoned Block Device models	17
4.2.1 Zoned Block Device models introduction	17
4.2.2 Host aware zoned block device model	19
4.2.3 Host managed zoned block device model	19
4.2.4 Domains and realms zoned block device model	21
4.2.4.1 Domains and realms zoned block device model overview	21
4.2.4.2 Zone domains	22
4.2.4.3 Zone domain 0	23
4.2.4.4 Zone domains other than zone domain 0	24
4.2.4.5 Zone activation	24
4.2.4.6 Realms	25
4.2.4.7 Realm boundary considerations	26
4.3 Zone attributes	27
4.3.1 Zone attributes summary	27
4.3.2 Zone Type zone attribute	28
4.3.3 Zone Condition zone attribute	29
4.3.4 WPointer zone attribute	29
4.3.5 RWP Recommended zone attribute	29
4.3.6 Non-Sequential Write Resources Active zone attribute	30
4.3.7 Predicted Unrecovered Errors Present zone attribute	30
4.4 Realm attributes	31
4.4.1 Realm attributes overview	31
4.4.2 Restrict Write Pointer Reset realm attribute	31
4.4.3 Restrict Zone Activate realm attribute	31
4.5 Zone type models	32
4.5.1 Zone type models overview	32
4.5.2 Conventional zone model	32
4.5.2.1 Conventional zone model overview	32
4.5.2.2 Write access pattern requirements for conventional zones	32

4.5.2.3 Read access pattern requirements for conventional zones	32
4.5.3 Write pointer zone models	32
4.5.3.1 Features common to all write pointer zones	32
4.5.3.1.1 Write pointer features	32
4.5.3.1.2 Resetting the write pointer	35
4.5.3.1.3 Open zone resources	35
4.5.3.1.4 Initialization pattern	35
4.5.3.1.5 Write access pattern requirements common to all write pointer zones	35
4.5.3.1.6 Read access pattern requirements common to all write pointer zones	36
4.5.3.2 Write pointer zone operations	37
4.5.3.2.1 Write pointer zone operations overview	37
4.5.3.2.2 Open zone operation	37
4.5.3.2.3 Close zone operation	38
4.5.3.2.4 Finish zone operation	38
4.5.3.2.5 Reset write pointer operation	38
4.5.3.2.6 Sequentialize zone operation	39
4.5.3.2.7 Zone activation operation	39
4.5.3.2.7.1 Zone activation operation overview	39
4.5.3.2.7.2 Verify activations operation	39
4.5.3.2.7.3 Change activations operation	40
4.5.3.2.8 Manage open zone resources operation	41
4.5.3.2.8.1 Manage open zone resources operation overview	41
4.5.3.2.8.2 Select a sequential write preferred zone	42
4.5.3.2.8.3 Select a sequential write required zone	42
4.5.3.2.9 Read operations, verify operations, and write operations	43
4.5.3.3 Sequential write preferred zone model	44
4.5.3.3.1 Sequential write preferred zone model overview	44
4.5.3.3.2 Write access pattern requirements for sequential write preferred zones	44
4.5.3.3.3 Read access pattern requirements for sequential write preferred zones	46
4.5.3.4 Sequential write required zone model	46
4.5.3.4.1 Sequential write required zone model overview	46
4.5.3.4.2 Write access pattern requirements for sequential write required zones	46
4.5.3.4.3 Read access pattern requirements for sequential write required zones	47
4.5.3.4.4 Opening Sequential Write Required zones	48
4.5.3.5 Sequential or before required zone model	48
4.5.3.5.1 Sequential or before required zone model overview	48
4.5.3.5.2 Write access pattern requirements for sequential or before required zones	48
4.5.3.5.3 Read access pattern requirements for sequential or before required zones	49
4.5.3.6 Zone condition state machine	50
4.5.3.6.1 Zone condition state machine overview	50
4.5.3.6.2 ZC1:Empty state	53
4.5.3.6.2.1 ZC1:Empty state overview	53
4.5.3.6.2.2 Transition ZC1:Empty to ZC2:Implicit_Open	53
4.5.3.6.2.3 Transition ZC1:Empty to ZC3:Explicit_Open	53
4.5.3.6.2.4 Transition ZC1:Empty to ZC6:Read_Only	53
4.5.3.6.2.5 Transition ZC1:Empty to ZC7:Offline	53
4.5.3.6.2.6 Transition ZC1:Empty to ZC8:Inactive	54
4.5.3.6.3 ZC2:Implicit_Open state	54
4.5.3.6.3.1 ZC2:Implicit_Open state overview	54
4.5.3.6.3.2 Transition ZC2:Implicit_Open to ZC1:Empty	54
4.5.3.6.3.3 Transition ZC2:Implicit_Open to ZC3:Explicit_Open	54
4.5.3.6.3.4 Transition ZC2:Implicit_Open to ZC4:Closed	54
4.5.3.6.3.5 Transition ZC2:Implicit_Open to ZC5:Full	55
4.5.3.6.3.6 Transition ZC2:Implicit_Open to ZC6:Read_Only	55
4.5.3.6.3.7 Transition ZC2:Implicit_Open to ZC7:Offline	55
4.5.3.6.4 ZC3:Explicit_Open state	55
4.5.3.6.4.1 ZC3:Explicit_Open state overview	55

4.5.3.6.4.2 Transition ZC3:Explicit_Open to ZC1:Empty	55
4.5.3.6.4.3 Transition ZC3:Explicit_Open to ZC4:Closed	56
4.5.3.6.4.4 Transition ZC3:Explicit_Open to ZC5:Full	56
4.5.3.6.4.5 Transition ZC3:Explicit_Open to ZC6:Read_Only	56
4.5.3.6.4.6 Transition ZC3:Explicit_Open to ZC7:Offline	56
4.5.3.6.5 ZC4:Closed state	56
4.5.3.6.5.1 ZC4:Closed state overview	56
4.5.3.6.5.2 Transition ZC4:Closed to ZC1:Empty	57
4.5.3.6.5.3 Transition ZC4:Closed to ZC2:Implicit_Open	57
4.5.3.6.5.4 Transition ZC4:Closed to ZC3:Explicit_Open	57
4.5.3.6.5.5 Transition ZC4:Closed to ZC6:Read_Only	57
4.5.3.6.5.6 Transition ZC4:Closed to ZC7:Offline	57
4.5.3.6.6 ZC5:Full state	57
4.5.3.6.6.1 ZC5:Full state overview	57
4.5.3.6.6.2 Transition ZC5:Full to ZC1:Empty	58
4.5.3.6.6.3 Transition ZC5:Full to ZC6:Read_Only	58
4.5.3.6.6.4 Transition ZC5:Full to ZC7:Offline	58
4.5.3.6.7 ZC6:Read_Only state	58
4.5.3.6.7.1 ZC6:Read_Only state overview	58
4.5.3.6.7.2 Transition ZC6:Read_Only to ZC7:Offline	58
4.5.3.6.8 ZC7:Offline state	59
4.5.3.6.9 ZC8:Inactive state	59
4.5.3.6.9.1 ZC8:Inactive state overview	59
4.5.3.6.9.2 Transition ZC8:Inactive to ZC1:Empty state	59
4.5.4 Gap zone model	59
4.6 Zoned block device extensions to block device model	60
4.6.1 Overview	60
4.6.2 Zoned block device internal resource management	60
4.6.3 Unexpected power removal	60
4.6.4 Media failure	61
4.7 Interactions involving mode parameter block descriptors	61
4.8 Capacity reporting and LBAs out of range	62
4.9 Constant zone starting LBA offsets	62
4.10 Format operations	63
4.11 Sanitize operations	64
4.12 Reservations	64
4.13 Caches	65
4.13.1 Caches overview	65
4.13.2 Write caching	66
4.13.3 Command interactions with caches	66
4.13.4 Write operation and write medium operation interactions with caches	66
4.13.5 Close zone and finish zone operation interactions with cache	66
4.14 Interactions with WRITE LONG commands	66
4.15 Interactions with storage element depopulation and restoration	67
4.15.1 Interactions with storage element depopulation and restoration operations that modify data	67
4.15.2 Storage element depopulation with zone modifications	67
4.15.2.1 Depopulation with zone modifications overview	67
4.15.2.2 Depopulation with zone modifications processing	67
4.15.2.3 Handling unrecoverable errors	68
4.15.2.3.1 Handling unrecoverable errors overview	68
4.15.2.3.2 Predicted unrecovered read errors in Conventional zones	69
4.15.2.3.3 Predicted unrecovered write errors in Conventional zones	69
4.15.2.4 Allowed commands during depopulation with zone modifications processing	70
4.15.2.5 Event handling actions	70
5 Commands for zoned block devices	71
5.1 Commands for zoned block devices overview	71

5.1.1 Summary of commands for zoned block devices	71
5.1.2 Zoned block device 16-byte CDB format with no data transfer	71
5.2 CLOSE ZONE command	74
5.3 FINISH ZONE command	76
5.4 OPEN ZONE command	78
5.5 REMOVE ELEMENT AND MODIFY ZONES command	79
5.6 REPORT REALMS command	81
5.6.1 REPORT REALMS command overview	81
5.6.2 REPORT REALMS parameter data	83
5.6.2.1 REPORT REALMS parameter data overview	83
5.6.2.2 Realm descriptor	85
5.6.2.2.1 Realm descriptor overview	85
5.6.2.2.2 Realm Start/End descriptor	86
5.7 REPORT ZONE DOMAINS command	87
5.7.1 REPORT ZONE DOMAINS command overview	87
5.7.2 REPORT ZONE DOMAINS parameter data	89
5.8 REPORT ZONES command	92
5.8.1 REPORT ZONES command overview	92
5.8.2 REPORT ZONES parameter data	94
5.9 RESET WRITE POINTER command	98
5.10 SEQUENTIALIZE ZONE command	100
5.11 ZONE ACTIVATE command	102
5.11.1 ZONE ACTIVATE command overview	102
5.11.2 Identifying the candidate zones to activate and the candidate zones to deactivate	103
5.11.3 ZONE ACTIVATE parameter data and ZONE QUERY parameter data	105
5.11.3.1 ZONE ACTIVATE parameter data and ZONE QUERY parameter data overview	105
5.11.3.2 Zone activation descriptors	107
5.12 ZONE QUERY command	109
6 Parameters for zoned block devices	110
6.1 Parameters for zoned block devices overview	110
6.2 Diagnostic parameters	110
6.3 Log parameters	111
6.3.1 Log parameters overview	111
6.3.2 Zoned Block Device Statistics log page	112
6.3.2.1 Zoned Block Device Statistics log page overview	112
6.3.2.2 Maximum Open Zones	114
6.3.2.3 Maximum Explicitly Open Zones	115
6.3.2.4 Maximum Implicitly Open Zones	116
6.3.2.5 Minimum Empty Zones	117
6.3.2.6 Maximum Non-sequential Zones	118
6.3.2.7 Zones Emptied	119
6.3.2.8 Suboptimal Write Commands	120
6.3.2.9 Commands Exceeding Optimal Limit	121
6.3.2.10 Failed Explicit Opens	122
6.3.2.11 Read Rule Violations	123
6.3.2.12 Write Rule Violations	124
6.3.2.13 Maximum Implicitly Open Sequential Or Before Required Zones	125
6.4 Mode parameters	126
6.4.1 Mode parameters overview	126
6.4.2 Zoned Block Device Control mode page	127
6.5 Vital product data (VPD) parameters	128
6.5.1 VPD parameters overview	128
6.5.2 Zoned Block Device Characteristics VPD page	129
Annex A (normative) ZBC Feature Sets	132
A.1 ZBC feature sets overview	132

A.2 Host Aware 2020 feature set	132
A.3 Host Managed 2020 feature set	134
A.4 Domains and Realms 2020 feature set	135
Annex B (informative) Application Client Considerations for Zoned Block Devices	137
B.1 Application client considerations for zoned block devices overview	137
B.2 Writing to write pointer zones	137
B.3 Open zone considerations	137
B.3.1 Open zone considerations overview	137
B.3.2 Explicitly opened zones and implicitly opened zones	138
B.3.3 Opening and closing zones	138
B.3.4 Finish zone operation considerations	139
B.4 Open zone resources considerations based on zone type	139
B.4.1 Sequential write preferred zones	139
B.4.2 Sequential write required zones	140
B.5 Partial failures	140
B.5.1 Partial failures overview	140
B.5.2 Sanitize considerations	140
Annex C (Informative) Bibliography	141

Tables

	Page
Table 1 – Numbering conventions	11
Table 2 – Comparison of decimal prefixes and binary prefixes	12
Table 3 – Example of ordering of bits and bytes within a data dword	13
Table 4 – Example of ordering of bits and bytes within an element dword	13
Table 5 – Zoned block device model concepts	15
Table 6 – Requirements of zoned block devices	18
Table 7 – Commands for host managed zoned block devices	20
Table 8 – Zone domain ID values	23
Table 9 – Summary of zone attributes	27
Table 10 – Zone Type zone attribute	28
Table 11 – Relationships between zone attributes	28
Table 12 – Zone Condition zone attribute	29
Table 13 – Summary of realm attributes	31
Table 14 – Summary of write pointer zone operations	37
Table 15 – Characteristics associated with zone state	51
Table 16 – READ CAPACITY (16) parameter data as modified for zoned block devices	62
Table 17 – RC BASIS field	62
Table 18 – ZBC-2 commands that are allowed in the presence of various reservations	65
Table 19 – Summary of commands that are unique to zoned block devices	71
Table 20 – Typical 16-byte zoned block device CDB format with no data transfer	72
Table 21 – CLOSE ZONE command	74
Table 22 – CLOSE ZONE command processing	75
Table 23 – FINISH ZONE command	76
Table 24 – FINISH ZONE command processing	77
Table 25 – OPEN ZONE command	78
Table 26 – OPEN ZONE command processing	79
Table 27 – REMOVE ELEMENT AND MODIFY ZONES command	79
Table 28 – REPORT REALMS command	81
Table 29 – REPORT REALMS REPORTING OPTIONS field	82
Table 30 – REPORT REALMS parameter data	83
Table 31 – Realm descriptor	85
Table 32 – REALM RESTRICTIONS field	86
Table 33 – Realm Start/End descriptor	86
Table 34 – REPORT ZONE DOMAINS command	87
Table 35 – REPORT ZONE DOMAINS REPORTING OPTIONS field	88
Table 36 – REPORT ZONE DOMAINS parameter data	89
Table 37 – Zone domain descriptor	90
Table 38 – REPORT ZONES command	92
Table 39 – REPORT ZONES REPORTING OPTIONS field	93
Table 40 – REPORT ZONES parameter data	94
Table 41 – SAME field description	95
Table 42 – Zone descriptor format	96
Table 43 – Zone descriptor ZONE TYPE field	96
Table 44 – Zone descriptor ZONE CONDITION field	97
Table 45 – RESET WRITE POINTER command	98
Table 46 – RESET WRITE POINTER command processing	99
Table 47 – SEQUENTIALIZE ZONE command	100
Table 48 – SEQUENTIALIZE ZONE command processing	101
Table 49 – ZONE ACTIVATE command	102
Table 50 – Selecting candidate zones to activate and deactivate with ALL bit set to zero	104
Table 51 – ZONE ACTIVATE parameter data and ZONE QUERY parameter data	105
Table 52 – Zone activation descriptor	108
Table 53 – ZONE QUERY command	109
Table 54 – Parameters for zoned block devices	110
Table 55 – Diagnostic page codes for host managed zoned block devices	110

Table 56 – Log page codes and subpage codes for host managed zoned block devices	111
Table 57 – Zoned Block Device Statistics log page parameter codes	112
Table 58 – Zoned Block Device Statistics log page	113
Table 59 – Maximum Open Zones log parameter	114
Table 60 – Maximum Explicitly Open Zones log parameter	115
Table 61 – Maximum Implicitly Open Zones log parameter	116
Table 62 – Minimum Empty Zones log parameter	117
Table 63 – Maximum Non-sequential Zones log parameter	118
Table 64 – Zones Emptied log parameter	119
Table 65 – Suboptimal Write Commands log parameter	120
Table 66 – Commands Exceeding Optimal Limit log parameter	121
Table 67 – Failed Explicit Opens log parameter	122
Table 68 – Read Rule Violations log parameter	123
Table 69 – Write Rule Violations log parameter	124
Table 70 – Maximum Implicitly Open Sequential Or Before Required Zones log parameter	125
Table 71 – Mode page codes and subpage codes for host managed zoned block devices	126
Table 72 – Zoned Block Device Control mode page	127
Table 73 – VPD page codes for zoned block devices	128
Table 74 – Zoned Block Device Characteristics VPD page	129
Table 75 – ZONED BLOCK DEVICE EXTENSION field	130
Table 76 – ZONE ALIGNMENT METHOD field	130
Table A.1 – Feature sets	132
Table A.2 – Commands mandatory for the Host Aware 2020 feature set	132
Table A.3 – Mode pages mandatory for the Host Aware 2020 feature set	133
Table A.4 – VPD pages mandatory for the Host Aware 2020 feature set	133
Table A.5 – Commands mandatory for the Host Managed 2020 feature set	134
Table A.6 – Mode pages mandatory for the Host Managed 2020 feature set	134
Table A.7 – VPD pages mandatory for the Host Managed 2020 feature set	134
Table A.8 – Commands mandatory for the Domains and Realms 2020 feature set	135
Table A.9 – Mode pages mandatory for the Domains And Realms 2020 feature set	135
Table A.10 – VPD pages mandatory for the Domains And Realms 2020 feature set	136

Figures

	Page
Figure 0 – SCSI document structure	xiv
Figure 1 – Example state diagram	14
Figure 2 – Zones in a zoned block device	17
Figure 3 – Zone domain	22
Figure 4 – Example of two zone domains	23
Figure 5 – Example of zone activation with the AAORB bit set to zero	24
Figure 6 – Realms model	25
Figure 7 – Example of conventional and shingled recording technologies using two zone domains	25
Figure 8 – Example of zone activation with the AAORB bit set to one	26
Figure 9 – Write pointer zone and write pointer after reset write pointer operation with no subsequent writes	33
Figure 10 – Write pointer zone and write pointer	33
Figure 11 – Write pointer zone example operations	34
Figure 12 – Example write command that starts at the write pointer	44
Figure 13 – Examples of write commands that do not start at the write pointer	45
Figure 14 – Zone condition state machine	52
Figure 15 – Example of paired sequential write required zones and gap zones	63

Foreword (This foreword is not part of American National Standard INCITS 550-2023.)

This purpose of this standard is to define the model and command set extensions to be used in conjunction with the SCSI Primary Command Set standard – 6 (SPC-6) and the SCSI Block Commands standard – 5 (SBC-5) to facilitate operation of zoned block devices.

Requests for interpretation, suggestions for improvement and addenda, or defect reports are welcome. They should be sent to the INCITS Secretariat, InterNational Committee for Information Technology Standards, Information Technology Industry Council, 700 K Street NW, Suite 600, Washington, DC 20001.

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

Laura Lindsay, Chair
Anish Karmarkar, Vice-Chair
Jennifer Garner, Secretary

<i>Organization Represented</i>	<i>Name of Representative</i>
Adobe Systems, Inc.	Scott Foshee
	David Sankel (Alt.)
AIM Global, Inc.	Steve Halliday
	Mary Lou Bosco (Alt.)
	Chuck Evanhoe (Alt.)
Amazon Web Services, Inc.	Tyler Messa
	Roland Blanding (Alt.)
	Rich Gardner (Alt.)
	Judd Hesselroth (Alt.)
Apple	Helene Workman
	Hicham Lozi (Alt.)
	Anna Weinberg (Alt.)
Dell, Inc.	Erin Bournival
Department of Commerce - NIST	Michael Hogan
	Lisa Carnahan (Alt.)
	Jessica Dickson (Alt.)
	Sal Francomacaro (Alt.)
	Suzanne Radell (Alt.)
Edison Design Group	John Spicer
	Daveed Vandevoorde (Alt.)
Farance, Inc.	Frank Farance
	Timothy Schoechle (Alt.)
Federal Bureau of Investigation (FBI)	Robert Brown
	Tara Alhariri (Alt.)
	Darrin Jones (Alt.)
	Kathleen Noyes (Alt.)
Futurewei Technologies, Inc.	Wael Diab
	Wilbert Adams (Alt.)
	Timothy Jeffries (Alt.)
GDCE, Inc.	Scott Payton
George Mason University	Duminda Wijesekera
Google	Catherine Nelson
	Ryan Wagner (Alt.)
	Lenora Zimmerman (Alt.)
GS1GO..... ..	Steven Keddle
	Edward Merrill (Alt.)
	Dan Mullen (Alt.)
HP, Inc.	Muhammad Ali
	Paul Jeran (Alt.)
IBM Corporation	Arnaud Le Hors
	Alexander Tarpinian (Alt.)

<i>Organization Represented</i>	<i>Name of Representative</i>
Immersion4 LLC	Serge Conesa
Intel Corporation.....	Dawn Blackbird (Alt.)
	Philip Wennblom
	Grace Wei (Alt.)
Microsoft Corporation	Laura Lindsay
	John Calhoon (Alt.)
NVidia Corporation	Timothy Costa
	Bryce Adelstein Leibach (Alt.)
Oracle.....	Anish Karmarkar
	Michael Kavanaugh (Alt.)
	Luke Kowalski (Alt.)
	Jan-Eike Michels (Alt.)
	Elaine Newton (Alt.)
Qualcomm, Inc.	Michael Atlass
	Mazen Chmaytelli (Alt.)
	Osok Song (Alt.)
Sandflow Consulting LLC	Pierre-Anthony Lemieux
United States Dept. of Defense.....	Vernita Harris
	Matthew Ivey (Alt.)
	Tejas Patel (Alt.)

INCITS Technical Committee T10 on SCSI Storage Interfaces, which developed and reviewed this standard, had the following members:

William Martin, Chair
Curtis Ballard, Vice-Chair
Curtis Stevens, Secretary
Frederick Knight, International Representative

<i>Organization Represented</i>	<i>Name of Representative</i>
Amphenol Corporation	Paul Coddington
	Brad Brubaker (Alt.)
	Zhineng Fan (Alt.)
	Adrian Green (Alt.)
	Donald Harper (Alt.)
	Michael Klempa (Alt.)
	Sam Kocsis (Alt.)
	Martin Li (Alt.)
	Chris Lyon (Alt.)
	Alex Persaud (Alt.)
	Michael Scholeno (Alt.)
	Hu Silver (Alt.)
	Michael Wingard (Alt.)
	Matt Wright (Alt.)
Anritsu Corporation	Tadanori Nishikobara
	Yuji Sekine (Alt.)
Broadcom, Inc.	Brad Besmer
	Srikiran Dravida (Alt.)
	John Gardner (Alt.)
	Jeffrey Gauvin (Alt.)
	Rick Kutcipal (Alt.)
	Bernhard Laschinsky (Alt.)
	David Peterson (Alt.)
	James Smart (Alt.)
	Jason Stuhlsatz (Alt.)
	Steven Wilson (Alt.)
Dell, Inc.	Kevin Marks
	David Black (Alt.)
	George Ericson (Alt.)
	Christopher Goonan (Alt.)
	Bill Lynn (Alt.)
	Gary Kotzur (Alt.)
	Bill Lynn (Alt.)
	Marlon Ramroopsingh (Alt.)

ISO/IEC 14776-346:2024(en)

<i>Organization Represented</i>	<i>Name of Representative</i>
Flextronics International (Taiwan), Ltd.	Roger Wan
Foxconn Interconnect Technology, Ltd. (FIT).....	Istvan BakroNagy (Alt.) Fred Fons Gary Hsieh (Alt.) Glenn Moore (Alt.) Miller Zhao (Alt.)
Futurewei Technologies, Inc.....	Victor Cao Timothy Jeffries (Alt.) Nelson Liao (Alt.)
Google	Radha Ramachandran Thieu Le (Alt.) Chris Sabol (Alt.) Bart Van Assche (Alt.)
Hewlett Packard Enterprise	Curtis Ballard Chris Cheng (Alt.) Robert Elliott (Alt.) Jeff Wolford (Alt.)
IBM Corporation	Kevin Butt
Intel Corporation	Mike Osborne (Alt.) Michael Hoard Phil Cayton (Alt.) Kim Malone (Alt.) Minesh Patel (Alt.) Eric Pickering (Alt.) Tzewen Wang (Alt.) Juntao Yuan (Alt.)
Keysight Technologies, Inc.....	Vincent Yew
Kioxia Corporation.....	John Geldman Cameron Brett (Alt.) Mark Carlson (Alt.) Don Harwood (Alt.) Koichi Nagai (Alt.) Paul Suhler (Alt.) Tatsuya Tanaka (Alt.)
LUXSHARE-ICT, Inc.	Scott Shuey Josue Castillo (Alt.) Jinhua Chen (Alt.) Pat Young (Alt.)
Marvell Semiconductor, Inc.	Craig Carlson Wei Liu (Alt.) James Walch (Alt.)
Meta Platforms, Inc.....	Ross Stenfort
Microchip Technology, Inc.....	Jermiah Tussey Sapna Devi (Alt.) Sanjay Goyal (Alt.) Adnan Jiwani (Alt.) Chethen K (Alt.) Keith Shaw (Alt.) Ariel Sibley (Alt.) Tim Symons (Alt.) Gregory Tabor (Alt.)
Molex, Inc.	Alexandra Haser Darian Schulz (Alt.) Scott Sommers (Alt.)
NetApp, Inc.....	Frederick Knight Chris Fore (Alt.) Jaimon George (Alt.) John Meneghini (Alt.)
Oracle	Martin Petersen Jon Allen (Alt.) Ioannis Psyllas (Alt.)
QNAP Systems, Inc.....	Ming-chih Chang CH Yang (Alt.)
Quantum Corporation	Darryl Torske Carsten Prigge (Alt.)

ISO/IEC 14776-346:2024(en)

<i>Organization Represented</i>	<i>Name of Representative</i>
Samsung Semiconductor, Inc. (SSI)	William Martin Judy Brock (Alt.) HeeChang Cho (Alt.) MiKyeong Kang (Alt.) Sung Lee (Alt.)
Seagate Technology	Gerald Houlder Umamahesh Allenki (Alt.) James Borden (Alt.) Andre Boyogueno (Alt.) Alvin Cox (Alt.) Ian Davies (Alt.) John Fleming (Alt.) Jim Hatfield (Alt.) Gerald Houlder (Alt.) Anil Keste (Alt.) Tony Kilwein (Alt.) Curtis Stevens (Alt.) Victor Taberski (Alt.) Judy Westby (Alt.)
Silicon Motion, Inc.	Amanda Huang Edward Hsieh (Alt.) Randy Hsu (Alt.) Sky Hsu (Alt.)
SK Hynix, Inc.	Taeho Lee Youngkyu Jeon (Alt.) Myron Loewen (Alt.) Jungki Noh (Alt.) Junho Son (Alt.)
Socionext	Masaru Suzuki Toru Iwata (Alt.) Hideyuki Kabuo (Alt.) Masanori Okinoi (Alt.)
TE Connectivity	Matt Schumacher Daniel Gorenc (Alt.) Tom Grzysiewicz (Alt.) Kyle Klinger (Alt.) Jeff Mason (Alt.) Joel Meyers (Alt.) Nathan Tracy (Alt.)
Teledyne LeCroy Corporation	Amit Bakshi Rakesh Kathiresan Natarajan (Alt.) Douglas Lee (Alt.)
Toshiba America Electronic Components, Inc.	Patrick Hery Johanna Hernandez (Alt.)
Unisys Corporation	Jeffrey Dremann Phil Shelton (Alt.)
Viavi Solutions, Inc.	Jason Rusch George Bullis (Alt.) Dominic Coupal (Alt.) Paul Gentieu (Alt.) Jason Mann (Alt.) Francisco Roy (Alt.) Yamini Shastry (Alt.)
VMware, Inc.	Neil MacLean Wenchao Cui (Alt.) Wenhua Liu (Alt.) Mike Panas (Alt.) Murali Rajagopal (Alt.) Ahmad Tawil (Alt.)

Organization Represented	Name of Representative
Western Digital Corporation	Dave Landsman
	Jorge Campello (Alt.)
	Xu Carter (Alt.)
	Marvin DeForest (Alt.)
	Kirill Dimitrov (Alt.)
	Michael Koffman (Alt.)
	Larry McMillian (Alt.)
	Chet Mercardo (Alt.)
	Nadesan Narenthiran (Alt.)
	Christopher Reed (Alt.)
	Yoni Shternhell (Alt.)
	Ralph Weber (Alt.)
	Eric Zornberg (Alt.)
<u>Emeritus</u>	
William Ham	
John Lohmeyer	

Information technology — Small computer system interface (SCSI) —

Part 346: Zoned Block Commands - 2 (ZBC-2)

1 Scope

This standard defines the model and command set extensions to facilitate operation of zoned block devices. The clauses in this standard, implemented in conjunction with the applicable clauses of SPC-6 and SBC-5, specify the standard command set for zoned block devices.

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.

T10/BSR INCITS 546, *SCSI Architecture Model - 6 (SAM-6)* (planned as ISO/IEC 14776-416)

T10/BSR INCITS 557, *SCSI / ATA Translation - 5 (SAT-5)* (under national consideration)

T10/BSR INCITS 566, *SCSI Primary Commands - 6 (SPC-6)* (planned as ISO/IEC 14776-456)

T10/BSR INCITS 571, *SCSI Block Commands - 5 (SBC-5)* (planned as ISO/IEC 14776-325)

3 Definitions, symbols, abbreviations, and conventions

3.1 Definitions

3.1.1 access pattern requirements

requirements, in addition to requirements in SBC-5, that apply to read operations, verify operations, and write operations performed by the device server in a zoned block device

3.1.2 activate

cause a write pointer zone to become an active zone

3.1.3 active zone

write pointer zone that has a Zone Condition other than INACTIVE

3.1.4 additional sense code

combination of the ADDITIONAL SENSE CODE field and the ADDITIONAL SENSE CODE QUALIFIER field in sense data

Note 1 to entry: See SPC-6.

3.1.5 application client

object that is the source of SCSI commands

Note 1 to entry: See SAM-6.

3.1.6 byte

8-bit construct

3.1.7 cache

temporary data storage area that is capable of containing a subset of the logical block data stored by the logical unit and is either volatile or non-volatile

Note 1 to entry: See 4.13.

3.1.8 close zone operation

device server process that makes zone resources non-volatile, releases open zone resources, and results in the Zone Condition becoming CLOSED or EMPTY

Note 1 to entry: See 4.5.3.2.3.

3.1.9 command

request describing a unit of work to be performed by a device server

Note 1 to entry: See SAM-6.

3.1.10 command descriptor block (CDB)

structure used to communicate commands from an application client to a device server

Note 1 to entry: See SPC-6.

3.1.11 conventional zone

zone that is not associated with a write pointer

Note 1 to entry: See 4.5.2.

3.1.12 deactivate

cause a write pointer zone to become an inactive zone

3.1.13 device server

object within a logical unit (see 3.1.28) that processes SCSI commands according to the rules of command management

Note 1 to entry: See SAM-6.

3.1.14 ending LBA

highest numbered LBA of a command

3.1.15 field

group of one or more contiguous bits, a part of a larger structure

Note 1 to entry: Examples of larger structures are a CDB (see 3.1.10) and sense data (see SPC-6).

3.1.16 finish zone operation

device server process that results in the Zone Condition becoming FULL

Note 1 to entry: See 4.5.3.2.4.

3.1.17 format operation

process by which a device server initializes the medium in a logical unit

Note 1 to entry: See 4.10.

3.1.18 gap zone

zone in which all of the LBAs are incapable of storing user data

Note 1 to entry: Gap zones enable the alignment of the starting LBAs of zones that are capable of storing user data (e.g., as described in 4.2.4.1 and 4.9).

Note 2 to entry: See 4.5.4.

3.1.19 hard reset

condition resulting from the events defined by SAM-6 during which the SCSI device performs the hard reset operations described in SAM-6, this standard, and other applicable command standards

3.1.20 host aware zoned block device

logical unit that implements the host aware zoned block device model

Note 1 to entry: See 4.2.2.

3.1.21 host managed zoned block device

logical unit that implements the host managed zoned block device model

Note 1 to entry: See 4.2.3.

3.1.22 inactive zone

write pointer zone that has a Zone Condition of INACTIVE

3.1.23 I_T nexus

relationship between a SCSI initiator port and a SCSI target port

Note 1 to entry: See SAM-6.

3.1.24 logical block

set of data bytes accessed and referenced as a unit

Note 1 to entry: See SBC-5.

3.1.25 logical block access command

command that requests access to one or more logical blocks that may require access to the medium

Note 1 to entry: See SBC-5.

3.1.26 logical block address (LBA)

value used to reference a logical block

Note 1 to entry: See SBC-5.

3.1.27 logical block data

user data and protection information, if any

3.1.28 logical unit

externally addressable entity within a SCSI target device (see 3.1.48) that implements a SCSI device model

Note 1 to entry: See SAM-6.

3.1.29 manage open zone resources operation

device server process that determines the availability of open zone resources

Note 1 to entry: See 4.5.3.2.8.

3.1.30 media

plural of medium

3.1.31 medium

material that is not cache on which data is stored (e.g., a magnetic disk)

3.1.32 non-sequential write operation

write operation performed on a write pointer zone for which the lowest LBA for that operation is not equal to the write pointer for that write pointer zone

Note 1 to entry: See 4.5.3.3.2.

3.1.33 non-volatile cache

cache that retains logical block data through any power cycle

3.1.34 open zone

zone with a Zone Condition of EXPLICITLY OPENED or IMPLICITLY OPENED

3.1.35 open zone operation

device server process that results in the Zone Condition becoming EXPLICITLY OPENED

3.1.36 open zone resources

zone resources that are available only to an open zone

Note 1 to entry: See 4.5.3.1.3.

3.1.37 persistent zone resources

zone resources that persist across all condition changes or SCSI events (see SAM-6) including a power on

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.

3.1.38 physical block

set of data bytes accessed as a unit by the device server (see SBC-5)

3.1.39 power cycle

sequence of power being removed followed by power being applied to a SCSI device

3.1.40 power on

condition resulting from the events defined by SAM-6 during which a SCSI device performs the power on operations described in SAM-6, this standard, and other applicable command standards

3.1.41 protection information

group of fields at the end of each logical block or at specified intervals within each logical block that contain a logical block guard, an application tag, and a reference tag

Note 1 to entry: See SBC-5.

3.1.42 read command

command that requests read operations

3.1.43 read operation

read operation as described in SBC-5 with the additional requirements described in this standard

3.1.44 realm

set of zone ranges, each in a different zone domain, and an indicated relationship among those zone ranges

Note 1 to entry: See 4.2.4.6.

3.1.45 realm attribute

named property associated with a realm

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

Note 2 to entry: See 4.4.

3.1.46 reset write pointer operation

device server process that results in the write pointer being set to the lowest LBA of the write pointer zone

Note 1 to entry: See 4.5.3.2.5.

3.1.47 sanitize operation

process by which a device server alters information on a logical unit such that recovery of previous logical block data from the cache and the medium is not possible

Note 1 to entry: See 4.11.

3.1.48 SCSI target device

SCSI device containing logical units and SCSI target ports that receives device service requests and task management requests for processing and sends device service responses and task management responses to SCSI initiator devices

Note 1 to entry: See SAM-6.

3.1.49 sense key

contents of the SENSE KEY field in the sense data

Note 1 to entry: See SPC-6.

3.1.50 sequential or before required zone

write pointer zone in which the device server allows write commands that specify a starting LBA that is less than or equal to that zone's write pointer

Note 1 to entry: See 4.5.3.5.

3.1.51 sequential write preferred zone

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

Note 1 to entry: See 4.5.3.3.

3.1.52 sequential write required zone

write pointer zone in which the device server requires that the lowest LBA for a write operation be the LBA indicated by the zone's write pointer

Note 1 to entry: See 4.5.3.4.

3.1.53 starting LBA

lowest numbered LBA of a command

3.1.54 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.

3.1.55 status

one byte of response information that contains a coded value defined in SAM-6, transferred from a device server to an application client upon completion of each command

Note 1 to entry: See SAM-6.

3.1.56 Trusted Computing Group

organization that develops and promotes open standards for hardware-enabled trusted computing and security technologies

Note 1 to entry: See <https://www.trustedcomputinggroup.org>.

3.1.57 unit attention condition

state that a logical unit (see 3.1.28) maintains while the logical unit has asynchronous status information to report to the SCSI initiator ports associated with one or more I_T nexuses (see 3.1.23)

Note 1 to entry: See SAM-6.

3.1.58 unmet prerequisite

requirement that, as a result of its not being met, prevents the successful performing of a zone activation operation (see 4.5.3.2.7)

Note 1 to entry: Unmet prerequisites are indicated by bits in the parameter data (see 5.11.3.1) returned by a command (e.g., the NOT EMPTY bit) as described in 4.5.3.2.7.2.

3.1.59 user data

data contained in logical blocks that is accessible by an application client and is neither protection information nor other information that may not be accessible to the application client

3.1.60 verify command

command that requests verify operations

3.1.61 verify operation

verify operation as described in SBC-5 with the additional requirements described in this standard

3.1.62 volatile cache

cache that does not retain logical block data between power cycles

3.1.63 write command

command that requests write operations

3.1.64 write operation

write operation as described in SBC-5 with the additional requirements described in this standard

3.1.65 write pointer

pointer to a logical block in a write pointer zone where the next write operation in that zone should start

Note 1 to entry: See 4.5.3.1.

3.1.66 write pointer zone

zone that has an associated write pointer

Note 1 to entry: See 4.5.3.

3.1.67 zone

one of a set of disjoint contiguous ranges of LBAs that comprise a logical unit

Note 1 to entry: See 4.5.

3.1.68 zone activation operation

operation that manages active zones and inactive zones

Note 1 to entry: See 4.5.3.2.7.

3.1.69 zone attribute

named property associated with a zone

Note 1 to entry: This standard assumes that all zone attributes (see 4.3) are zone resources regardless of how the device server maintains the defined information.

Note 2 to entry: In the absence of other conditions or actions (e.g., actions by the application client), zone attributes do not affect the operation of the zone condition state machine.

Note 3 to entry: See 4.3.

3.1.70 zone domain

zone range in which all of the zones are able to be active zones at the same time

Note 1 to entry: See 4.2.4.2.

3.1.71 zone ID

lowest LBA of a zone

3.1.72 zone range

contiguous set of zones specified by a zone ID and a non-zero integer number of zones

3.1.73 zone resources

device server resources that are used by a zone

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

3.1.74 zoned block device

host aware zoned block device, host managed zoned block device, or domains and realms zoned block device

Note 1 to entry: See 4.2.

3.1.75 zoned maximum address

maximum LBA as indicated by the MAXIMUM LBA field (see 5.8.2) returned by the REPORT ZONES command

Note 1 to entry: See 4.8.

3.2 Symbols and abbreviations

3.2.1 Abbreviations

Abbreviations used in this standard:

Abbreviation	Meaning
CDB	command descriptor block
LBA	logical block address
LSB	least significant bit
MSB	most significant bit
n/a	not applicable
SAM-6	SCSI Architecture Model - 6 (see clause 2)
SAT-5	SCSI / ATA Translation - 5 (see clause 2)
SBC-5	SCSI Block Commands - 5 (see clause 2)
SCSI	Small Computer System Interface family of standards
SPC-6	SCSI Primary Commands - 6 (see clause 2)
TCG	Trusted Computing Group (see 3.1.56)
VPD	Vital product data

3.2.2 Mathematical operators

Mathematical operators used in this standard:

Mathematical Operators	Meaning
$\wedge$ or XOR	exclusive logical OR
$\times$	multiplication
$/$	division
$\pm$	plus or minus
$+$	add
$-$	subtract
AND	logical AND
$<$ or LT	less than
$\leq$ or LE	less than or equal to
$=$ or EQ	equal
$\neq$ or NE	not equal
$>$ or GT	greater than
$\geq$ or GE	greater than or equal to

3.3 Keywords

3.3.1 invalid

keyword used to describe an illegal or unsupported bit, byte, word, field, or code value

Note 1 to entry: Receipt by a device server of an invalid bit, byte, word, field, or code value shall be reported as error.

3.3.2 mandatory

keyword indicating an item that is required to be implemented as defined in this standard

3.3.3 may

keyword that indicates flexibility of choice with no implied preference

3.3.4 may or may not

keyword that indicates flexibility of choice with no implied preference

Note 1 to entry: Significant uses of "may or may not" occur in descriptions where attention is being drawn to the "may not" case.

3.3.5 obsolete

keyword indicating that an item was defined in prior SCSI standards but has been removed from this standard

3.3.6 option, optional

keywords that describe features that are not required to be implemented by this standard

Note 1 to entry: If any optional feature defined by this standard is implemented, then it shall be implemented as defined in this standard.

3.3.7 prohibited

keyword used to describe a feature, function, or coded value that is defined in a non-SCSI standard (i.e., a standard that is not a member of the SCSI family of standards) to which this standard makes a normative reference where the use of said feature, function, or coded value is not allowed for implementations of this standard

3.3.8 reserved

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

Note 1 to entry: A reserved bit, byte, word, or field shall be set to zero, or in accordance with a future extension to this standard.

Note 2 to entry: Recipients are not required to check reserved bits, bytes, words, or fields for zero values.

Note 3 to entry: Receipt of reserved code values in defined fields shall be reported as an error.

3.3.9 restricted

keyword referring to bits, bytes, words, and fields that are set aside for other identified standardization purposes

Note 1 to entry: A restricted bit, byte, word, or field shall be treated as a reserved bit, byte, word, or field in the context where the restricted designation appears.

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

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 SCSI device vendor and may be used differently in various implementations.

3.4 Editorial conventions

Certain words and terms used in this standard have a specific meaning beyond the normal English meaning. These words and terms are defined either in the glossary or in the text where they first appear.

Upper case is used when referring to the name of a numeric value defined in this specification or a formal attribute possessed by an entity. When necessary for clarity, names of objects, procedure calls, arguments or discrete states are capitalized or set in bold type. Names of fields are identified using small capital letters (e.g., OPERATION CODE field).

Quantities having a defined numeric value are identified by large capital letters (e.g., CHECK CONDITION). Quantities having a discrete but unspecified value are identified using small capital letters (e.g., OFFLINE is a value of the Zone Condition zone attribute). Such quantities are associated with an event or indication whose observable behavior or value is specific to a given implementation standard.

Lists sequenced by lowercase or uppercase letters show no ordering relationship between the listed items.

EXAMPLE 1 - The following list shows no relationship between the named items:

- a) red (i.e., one of the following colors):
 - A) crimson; or
 - B) amber;
- b) blue; or
- c) green.

Lists sequenced by numbers show an ordering relationship between the listed items.

EXAMPLE 2 - The following list shows an ordered relationship between the named items:

- 1) top;
- 2) middle; and
- 3) bottom.

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).

If a conflict arises between text, tables, or figures, the order of precedence to resolve the conflicts is:

- 1) text;
- 2) tables; and
- 3) figures.

Not all tables or figures are fully described in the text.

Notes and examples do not constitute any requirements.

Notes are numbered consecutively throughout this standard.

3.5 Numeric and character conventions

3.5.1 Numeric conventions

A binary number is represented in this standard by any sequence of digits comprised of only the Arabic numerals 0 and 1 immediately followed by a lower-case b (e.g., 0101b). Underscores or spaces may be included 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 comprised of only the Arabic numerals 0 to 9 and/or the upper-case English letters A to F immediately followed by a lower-case h (e.g., FA23h). Underscores or spaces may be included in hexadecimal number representations to increase readability.

or delineate field boundaries (e.g., 3456FDCA 84BD5E7Ah, 3456FDCA_84BD5E7Ah, B FD8C FA23h, or B_FD8C_FA23h).

A decimal number is represented in this standard by any sequence of digits comprised of only the Arabic numerals 0 to 9 not immediately followed by a lower-case b or lower-case h (e.g., 25).

Variables (i.e., alphanumeric names that represent values in computations and other statements) are represented in the same San-serif font as other information in this standard.

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;
- c) the thousands separator is used in both the integer portion and the fraction portion of a number; and
- d) the decimal representation for a year is 1999 not 1 999.

Table 1 shows some examples of decimal numbers using various 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 $12\ 1/7$).

A range of numeric values is represented in this standard in the form “a to z”, where a is the first value included in the range, all values between a and z are included in the range, and z is the last value included in the range (e.g., the representation “0h to 3h” includes the values 0h, 1h, 2h, and 3h).

3.5.2 Units of measure

This standard represents values using both decimal units of measure and binary units of measure. Values are represented by the following formats:

- a) for values based on decimal units of measure:
 - 1) numerical value (e.g., 100);
 - 2) space; and
 - 3) prefix symbol and unit:
 - 1) decimal prefix symbol (e.g., M) (see table 2); and
 - 2) unit abbreviation (e.g., B);

and
- b) for values based on binary units of measure:
 - 1) numerical value (e.g., 1 024);
 - 2) space; and
 - 3) prefix symbol and unit:
 - 1) binary prefix symbol (e.g., Gi) (see table 2); and
 - 2) unit abbreviation (e.g., b).

Table 2 compares the prefix, symbols, and power of the binary and decimal units.

Table 2 – Comparison of decimal prefixes and binary prefixes

Decimal			Binary		
Prefix name	Prefix symbol	Power (base-10)	Prefix name	Prefix symbol	Power (base-2)
kilo	k	10^3	kibi	Ki	2^{10}
mega	M	10^6	mebi	Mi	2^{20}
giga	G	10^9	gibi	Gi	2^{30}
tera	T	10^{12}	tebi	Ti	2^{40}
peta	P	10^{15}	pebi	Pi	2^{50}
exa	E	10^{18}	exbi	Ei	2^{60}
zetta	Z	10^{21}	zebi	Zi	2^{70}
yotta	Y	10^{24}	yobi	Yi	2^{80}

3.6 Bit and byte ordering

In this standard, data structures may be defined by a table. A table defines a complete ordering of elements (i.e., bits, bytes, fields, and dwords) within the structure. The ordering of elements within a table does not in itself constrain the order of storage or transmission of the data structure, but in combination with other normative text in this standard, may constrain the order of storage or transmission of the structure.

In a table, any element that is presented in a row above another element in a lower row is more significant than the lower element, and any element presented to the left of another element in the same row is more significant than the element to the right.

If a table shows bit numbering (see table 3), the least significant bit (LSB) is numbered 0 and each more significant bit has the next greater number than the immediately less significant bit. If a table shows numbering of bytes or characters (see table 4), the most significant byte or character is represented at the lowest number and each less significant byte or character has the next greater number than the immediately more significant byte.

In a field in a table consisting of more than one bit that contains a single value (e.g., a number), the least significant bit (LSB) is shown on the right and the most significant bit (MSB) is shown on the left (e.g., in a byte, bit 7 is the MSB and is shown on the left, bit 0 is the LSB and is shown on the right). The MSB and LSB are not labeled if the field consists of eight or fewer bits. The MSB and LSB are labeled if the field consists of more than eight bits and has no internal structure defined.

In a field in a table consisting of more than one byte that contains multiple fields each with their own values (e.g., a descriptor), there is no MSB and LSB of the field itself and thus there are no MSB and LSB labels. Each individual field has an MSB and LSB, but they are not labeled.

In a field containing a text string (e.g., ASCII or UTF-8), only the MSB of the first character and the LSB of the last character are labeled.

Multiple byte fields are represented with only two rows, with the non-sequentially increasing byte number denoting the presence of additional bytes.

A data dword consists of 32 bits. Table 3 shows a data dword containing a single value, where the MSB is on the upper left in bit 31 and the LSB is on the lower right in bit 0.

Table 3 – Example of ordering of bits and bytes within a data dword

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB) Bit 31	Bit 30	Bit 29	Bit 28	Bit 27	Bit 26	Bit 25	Bit 24
1	Bit 23	Bit 22	Bit 21	Bit 20	Bit 19	Bit 18	Bit 17	Bit 16
2	Bit 15	Bit 14	Bit 13	Bit 12	Bit 11	Bit 10	Bit 9	Bit 8
3	Bit 7	Bit 6	Bit 5	Bit 4	Bit 3	Bit 2	Bit 1	Bit 0 (LSB)
Note - The Bit x labels in the individual table cells are for reference only and should not appear within tables that use this element format.								

Table 4 shows a data dword containing four one-byte fields, where byte 0 (the first byte) is on the left and byte 3 (the fourth byte) is on the right. Each byte has an MSB on the left and an LSB on the right.

Table 4 – Example of ordering of bits and bytes within an element dword

Bit Byte	7	6	5	4	3	2	1	0
0	First byte							
	Bit 7 (MSB)	Bit 6	Bit 5	Bit 4	Bit 3	Bit 2	Bit 1	Bit 0 (LSB)
1	Second byte							
	Bit 7 (MSB)	Bit 6	Bit 5	Bit 4	Bit 3	Bit 2	Bit 1	Bit 0 (LSB)
2	Third byte							
	Bit 7 (MSB)	Bit 6	Bit 5	Bit 4	Bit 3	Bit 2	Bit 1	Bit 0 (LSB)
3	Fourth byte							
	Bit 7 (MSB)	Bit 6	Bit 5	Bit 4	Bit 3	Bit 2	Bit 1	Bit 0 (LSB)
Note - The Bit x labels in the individual table cells and the xx byte labels in the individual bytes are for reference only and should not appear within tables that use these element formats. In this example the MSB and LSB labels are for reference only, however, they may appear in multi-byte fields as described in this subclause (i.e., 3.6).								

3.7 Notation for state diagrams

All state diagrams use the notation shown in figure 1.




Figure 1 – Example state diagram

Each state is identified by a state designator and a state name. The state designator (e.g., B1) is unique among all state machines in this standard. The state name (e.g., State 1) is a brief description of the primary action taken during the state, and the same state name may be used by other state machines. Actions taken while in each state are described in the state description text.

Transitions between states are shown with solid lines, with an arrow pointing to the destination state.

The conditions and actions are described fully in the transition description text.

Transitions between states are instantaneous.

Upon entry into a state, all actions to be processed in that state are processed. If a state is re-entered from itself, all actions to be processed in the state are processed again. A state may be entered and exited in zero time if the conditions for exiting the state are valid upon entry into the state.

4 Zoned Block Device Model

4.1 Zoned Block Device model overview

4.1.1 Established SCSI concepts

Table 5 lists established SCSI concepts that apply to zoned block devices.

Table 5 – Zoned block device model concepts (part 1 of 2)

Concept ^a	Reference
Zoned block device models	4.2
Zone attributes	4.3
Realm attributes	4.4
Zone type models	4.5
Error reporting	4.5.3 and SBC-5
Zoned block device extensions to block device model	4.6
Capacity reporting and LBAs out of range	4.8
Constant zone starting LBA values	4.9
Format operations	4.10 and SBC-5
Sanitize operations	4.11 and SBC-5
Reservations	4.12, SPC-6, and SBC-5
Caches	4.13
Pseudo unrecovered errors	4.14 and SBC-5
Storage element depopulation and restoration operations that modify data	4.15.1 and SBC-5
Storage element depopulation with zone modifications	4.15.2 and SBC-5
Background scan operations	SBC-5
Deferred microcode activation	SBC-5
Grouping function	SBC-5
Implicit head of queue command processing	SBC-5
Initialization	SBC-5
Media examples	SBC-5
Medium defects	SBC-5
^a Except for host managed zoned block devices, zoned block devices may support any concept or model defined in SBC-5. Host managed zoned block devices should not support SBC-5 model subclauses not listed in this table.	

Table 5 – Zoned block device model concepts (part 2 of 2)

Concept ^a	Reference
Logical blocks	SBC-5
Physical blocks	SBC-5
Protection information model	SBC-5
Ready state	SBC-5
START STOP UNIT and power conditions	SBC-5
Write and unmap failures	SBC-5
Write protection	SBC-5
^a Except for host managed zoned block devices, zoned block devices may support any concept or model defined in SBC-5. Host managed zoned block devices should not support SBC-5 model subclauses not listed in this table.	

4.1.2 Peripheral device type and supported commands

Zoned block devices indicate their primary model using the PERIPHERAL DEVICE TYPE field in the standard INQUIRY data (see SPC-6). Peripheral device type of:

- a) 14h indicates a host managed zoned block device (see 4.2.3); and
- b) 00h (i.e., direct access block device) indicates:
 - A) a host aware zoned block device (see 4.2.2); or
 - B) a domains and realms zoned block device (see 4.2.4).

Application clients may use the contents of the PERIPHERAL DEVICE TYPE field as an indication of what commands may be supported by the device server, as follows:

- a) 00h indicates that command support requirements are defined in SBC-5 and this standard (see table 6); and
- b) 14h indicates that command support requirements are defined in this standard (see table 7).

Additional information may be available in the ZONED BLOCK DEVICE EXTENSION field in the Zoned Block Device Characteristics VPD page (see 6.5.2).

For commands that are unique to zoned block devices (see table 19), support requirements are shown in table 6.

4.2 Zoned Block Device models

4.2.1 Zoned Block Device models introduction

Zoned block device models described in this standard are as follows:

- a) host aware zoned block device (see 4.2.2);
- b) host managed zoned block device (see 4.2.3); and
- c) domains and realms zoned block device (see 4.2.4).

Zoned block devices are accessed using LBAs. The LBAs are divided into ranges called zones (see 4.5). The entire capacity of a zoned block device is organized into a set of logically contiguous, non-overlapping zones. Figure 2 shows a zoned block 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.

Zone	0		1	...	n-1	
LBA	0	1	...			m-1

Figure 2 – Zones in a zoned block device

Each zone of a zoned block device has an associated zone type (see 4.5). Zoned block devices support the REPORT ZONES command (see 5.8) for reporting the zone type and other information for each zone.

Table 6 defines the characteristics of zoned block devices defined in this standard.

Table 6 – Requirements of zoned block devices

Characteristic	Reference	Host aware zoned block device	Host managed zoned block device	Domains and realms zoned block device
Command support	4.1.2	See SPC-6 and SBC-5	See table 7	See SPC-6 and SBC-5
PERIPHERAL DEVICE TYPE field value	4.1.2	00h	14h	00h
ZONED BLOCK DEVICE EXTENSION field value in the Zoned Block Device Characteristics VPD page	6.5.2	1h	0h	2h
Conventional zone	4.5.2	Optional	Optional	Not supported
Sequential write preferred zone	4.5.3.3	Mandatory	Not supported	Optional ^a
Sequential write required zone	4.5.3.4	Not supported	Mandatory	Optional ^a
Sequential or before required zone	4.5.3.5	Not supported	Not supported	Mandatory
Gap zone	4.5.4	Not supported	see 4.2.3	Optional
CLOSE ZONE command	5.2	Mandatory	Mandatory	Mandatory
FINISH ZONE command	5.3	Mandatory	Mandatory	Mandatory
OPEN ZONE command	5.4	Mandatory	Mandatory	Mandatory
REMOVE ELEMENT AND MODIFY ZONES command	5.5	Optional	Optional	Optional
REPORT REALMS command	5.6	Not supported	Not supported	Optional
REPORT ZONE DOMAINS command	5.7	Not supported	Not supported	Mandatory
REPORT ZONES command	5.8	Mandatory	Mandatory	Mandatory
RESET WRITE POINTER command	5.9	Mandatory	Mandatory	Mandatory
SEQUENTIALIZE ZONE command	5.10	Mandatory	Not supported	Optional ^b
ZONE ACTIVATE command	5.11	Not supported	Not supported	Mandatory
ZONE QUERY command	5.12	Not supported	Not supported	Mandatory
^a The device server shall support one or more of the following types of zones: a) sequential write preferred zones; or b) sequential write required zones. ^b If the device server supports sequential write preferred zones (see 4.5.3.3), then the SEQUENTIALIZE ZONE command (see 5.10) is mandatory.				

4.2.2 Host aware zoned block device model

A host aware zoned block device reports:

- a) the PERIPHERAL DEVICE TYPE field set to 00h (i.e., direct access block device) (see SPC-6); and
- b) the ZONED BLOCK DEVICE EXTENSION field set to 1h in the Zoned Block Device Characteristics VPD page (see 6.5.2).

Host aware zoned block devices:

- a) may support conventional zones (see 4.5.2);
- b) shall support one or more sequential write preferred zones (see 4.5.3.3);
- c) shall not support sequential write required zones (see 4.5.3.4);
- d) shall not support sequential or before required zones (see 4.5.3.5); and
- e) shall not support gap zones (see 4.5.4).

Capacity reporting requirements are defined in 4.8.

In addition to primary commands (see SPC-6) and commands for direct access block devices (see SBC-5), host aware zoned block devices shall support the commands specified in table 6.

4.2.3 Host managed zoned block device model

A host managed zoned block device reports:

- a) the PERIPHERAL DEVICE TYPE field set to 14h (i.e., host managed zoned block device) (see SPC-6); and
- b) the ZONED BLOCK DEVICE EXTENSION field set to 0h in the Zoned Block Device Characteristics VPD page (see 6.5.2).

Host managed zoned block devices:

- a) may support conventional zones (see 4.5.2);
- b) shall not support sequential write preferred zones (see 4.5.3.3);
- c) shall support one or more sequential write required zones (see 4.5.3.4); and
- d) shall not support sequential or before required zones (see 4.5.3.5).

If the ZONE ALIGNMENT METHOD field (see 6.5.2) is:

- a) set to 8h, then the host managed zone block device may support gap zones (see 4.5.4) as described in 4.9; and
- b) not set to 8h, then the host managed zone block device shall not support gap zones.

Capacity reporting requirements are defined in 4.8.

Table 7 lists the commands for host managed zoned block devices.

Table 7 – Commands for host managed zoned block devices (part 1 of 2)

Command	Operation code	Type	LBACT	Reference
ATA PASS-THROUGH (12)	A1h	O	n/a	SAT-5
ATA PASS-THROUGH (16)	85h	O	n/a	SAT-5
CLOSE ZONE	94h/01h	M	Z	5.2
FINISH ZONE	94h/02h	M	Z	5.3
FORMAT UNIT	04h	O	Z	SBC-5 and 4.10
FORMAT WITH PRESET	38h	O	Z	SBC-5 and 4.10
INQUIRY	12h	M	n/a	SPC-6
LOG SELECT	4Ch	O	n/a	SPC-6
LOG SENSE	4Dh	M	n/a	SPC-6
MODE SELECT (10)	55h	M	n/a	SPC-6
MODE SENSE (10)	5Ah	M	n/a	SPC-6
OPEN ZONE	94h/03h	M	Z	5.4
PERSISTENT RESERVE IN	5Eh	O	n/a	SPC-6
PERSISTENT RESERVE OUT	5Fh	O	n/a	SPC-6
READ (16)	88h	M	R	SBC-5
READ BUFFER (10)	3C	O	n/a	SPC-6
READ BUFFER (16)	9Bh	O	n/a	SPC-6
READ CAPACITY (16)	9Eh/10h	M	n/a	SBC-5 and 4.8
READ DEFECT DATA (12)	B7h	O	n/a	SBC-5
REMOVE ELEMENT AND MODIFY ZONES	9Eh/1Ah	O	Z	5.5
REMOVE ELEMENT AND TRUNCATE	9Eh/18h	O	Z	SBC-5
REPORT LUNS	A0h	M	n/a	SPC-6
REPORT SUPPORTED OPERATION CODES	A3h/0Ch	M	n/a	SPC-6
Key: O = optional M = mandatory R = read command U = unmap command V = verify command W = write command Z = other command LBACT = logical block access command type (see SBC-5)				

Table 7 – Commands for host managed zoned block devices (part 2 of 2)

Command	Operation code	Type	LBACT	Reference
REPORT SUPPORTED TASK MANAGEMENT FUNCTIONS	A3h/0Dh	M	n/a	SPC-6
REPORT TIMESTAMP	A3/0Fh	O	n/a	SPC-6
REPORT ZONES	95h/00h	M	Z	5.8
REQUEST SENSE	03h	M	n/a	SPC-6
RESET WRITE POINTER	94h/04h	M	Z	5.9
RESTORE ELEMENT AND REBUILD	9Eh/19h	O	Z	SBC-5
SANITIZE	48h	O	Z	SBC-5 and 4.11
SECURITY PROTOCOL IN	A2h	O	n/a	SPC-6
SECURITY PROTOCOL OUT	B5h	O	n/a	SPC-6
SEND DIAGNOSTIC	1Dh	O	n/a	SPC-6
SET TIMESTAMP	A4h/0Fh	O	n/a	SPC-6
START STOP UNIT	1Bh	M	n/a	SBC-5
SYNCHRONIZE CACHE (16)	91h	M	W	SBC-5
TEST UNIT READY	00h	M	n/a	SPC-6
VERIFY (16)	8Fh	O	V, W	SBC-5
WRITE (16)	8Ah	M	W	SBC-5
WRITE BUFFER	3Bh	O	n/a	SPC-6
WRITE LONG (16)	9Fh/11h	O	Z	SBC-5 and 4.14
WRITE SAME (16)	93h	M	U, W	SBC-5
Key: O = optional M = mandatory R = read command U = unmap command V = verify command W = write command Z = other command LBACT = logical block access command type (see SBC-5)				

4.2.4 Domains and realms zoned block device model

4.2.4.1 Domains and realms zoned block device model overview

A domains and realms zoned block device reports:

- the PERIPHERAL DEVICE TYPE field set to 00h (i.e., direct access block device) (see SPC-6); and
- the ZONED BLOCK DEVICE EXTENSION field set to 2h in the Zoned Block Device Characteristics VPD page (see 6.5.2).

Domains and realms zoned block devices:

- a) shall not support conventional zones (see 4.5.2);
- b) shall support one or more of the following types of zones:
 - A) sequential write preferred zones (see 4.5.3.3); or
 - B) sequential write required zones (see 4.5.3.4);
- c) shall support one or more sequential or before required zones (see 4.5.3.5); and
- d) may support gap zones (see 4.5.4).

Capacity reporting requirements are defined in 4.8.

In addition to primary commands (see SPC-6) and commands for direct access block devices (see SBC-5), domains and realms zoned block devices shall support the commands specified in table 6.

Zone domains and realms enhance the basic zoned block device model (see 4.2.1) to define a mechanism that allows an application client to manage contention for certain SCSI device resources in a way that is independent of many details about those SCSI device resources. Two or more zone domains (see 4.2.4.2) represent the SCSI device resources for which an application client is able to manage contention.

The SCSI device resources for which contention management is provided include, but are not limited to, the ability to store logical block data. The contention for SCSI device resources is managed by having:

- a) some zones be active zones (i.e., able to use the SCSI device resource in contention); and
- b) the remaining zones be inactive zones (i.e. unable to use the SCSI device resource in contention).

The REPORT ZONES command (see 5.8), REPORT REALMS command (see 5.6), REPORT ZONE DOMAINS command (see 5.7), and the ZONE QUERY command (see 5.12) provides current indications of what zone domains and zone types are available for the application client to choose to make active zones.

The ZONE ACTIVATE command (see 5.11) may activate a zone range in one zone domain and may deactivate a corresponding zone range in one or more different zone domains.

Commands that specify LBAs (e.g., read commands and write commands) in inactive zones are processed as described in 4.5.3.1.5 and 4.5.3.1.6.

4.2.4.2 Zone domains

A zone domain (see figure 3) is a zone range that is identified by a zone domain ID (see table 8). A zone domain ID is assigned by the device server starting at zero and incrementing by one. The zone range that constitutes the Zone domain shown in figure 3:

- 1) starts with Zone t, followed by Zone t+1, Zone t+2, and Zone t+3;
- 2) continues for zero or more zones; and
- 3) ends with Zone x-3, followed by Zone x-2, Zone x-1 and Zone x.

Zone t is the first zone in the Zone domain shown in figure 3 and the lowest LBA in Zone t is the lowest LBA of the Zone domain. Zone x is the last zone in the Zone domain shown in figure 3 and the highest LBA in Zone x is the highest LBA of the Zone domain.




Figure 3 – Zone domain

Table 8 – Zone domain ID values

Code	Description
00h to F7h	Zone domain ID values
F8h to FEh	Reserved
FFh	No zone domain ID specified

All the zones in a single zone domain shall be capable of being active zones at the same time.

Each zone that is capable of being activated shall be included in one and only one zone domain.

Two different zone domains may have:

- a) different numbers of zones; and
- b) different numbers of LBAs.

Figure 4 shows an example of a device server that supports two zone domains with:

- 1) Zone domain 0:
 - A) starting with Zone a and ending with Zone b; and
 - B) containing:
 - a) LBA 0 as the lowest LBA in Zone domain 0; and
 - b) LBA m as the highest LBA in Zone domain 0;
- 2) zero or more gap zones (see 4.5.4) that separate Zone domain 0 from Zone domain 1, with
 - A) LBA m+1 as the lowest LBA of the lowest gap zone, if any; and
 - B) LBA p-1 as the highest LBA of the highest gap zone, if any;
 and
- 3) Zone domain 1:
 - A) starting with Zone y and ending with Zone z; and
 - B) containing:
 - a) LBA p as the lowest LBA in Zone domain 1; and
 - b) LBA r as the highest LBA in Zone domain 1.




Figure 4 – Example of two zone domains

4.2.4.3 Zone domain 0

A domains and realms zoned block device shall have one zone domain with LBA 0 as its lowest numbered LBA. That zone domain is identified as zone domain 0.

As described in 4.8, the highest LBA in zone domain 0 establishes the value returned in the RETURNED LOGICAL BLOCK ADDRESS field in the READ CAPACITY (16) parameter data.

At the time of manufacture, a domains and realms zoned block device shall have a zone domain 0 that contains only sequential or before required zones (see 4.5.3.5) that have a Zone Condition of FULL.

4.2.4.4 Zone domains other than zone domain 0

For zone domains other than zone domain 0, the device server shall support accesses (e.g., reads and writes) that specify LBAs that are greater than the contents of the RETURNED LOGICAL BLOCK ADDRESS field in the READ CAPACITY (16) parameter data (see SBC-5 and 4.8).

The LBAs in the zones in zone domains other than zone domain 0 are:

- a) greater than the value in the RETURNED LOGICAL BLOCK ADDRESS field; and
- b) less than or equal to the zoned maximum address.

4.2.4.5 Zone activation

A successful ZONE ACTIVATE command (see 5.11) that specifies a zone range to be activated activates one or more zones and deactivates zero or more zones. A successful ZONE ACTIVATE command that specifies a zone range to be deactivated activates zero or more zones and deactivates one or more zones.

The smallest zone range processed by a ZONE ACTIVATE command or a ZONE QUERY command (see 5.12) is:

- a) one zone, if the AAORB bit (see 6.5.2) is set to zero (see example in figure 5); and
- b) all the zones in one zone realm, if the AAORB bit is set to one (see example in figure 8).

Figure 5 shows an example of the before and after conditions for a ZONE ACTIVATE command that is processed with the AAORB bit set to zero.




Figure 5 – Example of zone activation with the AAORB bit set to zero

4.2.4.6 Realms

Combined with zone domains (see 4.2.4.2), realms form a matrix (see figure 6) that provides applications clients with a set of choices for which zones are active. In figure 6, zone domain 0 contains LBA 0 (see 4.2.4.3) and each realm contains more than one zone. Figure 3 describes how zones are contained within zone domains.

Realm 0	Realm 1	Realm 2	...	Realm r
	Zone Domain 0			
			...	
	Zone Domain 1			
		...	...	
	Zone Domain z			

Figure 6 – Realms model

Each realm shall contain at least one zone.

Realms associate characteristics with:

- a) zone domains (e.g., zone type, recording technology, data transfer performance); and
- b) realms (e.g., location on the media, data storage density, storage density performance).

EXAMPLE - Figure 7 shows a configuration in which one zone domain uses conventional recording in sequential or before required zones (see 4.5.3.5) and one zone domain uses shingled magnetic recording in sequential write required zones (see 4.5.3.4). In this configuration, the domains and realms zoned block device allows a host to modify which realms store logical block data without the full effects of shingled magnetic recording and which realms store logical block data using shingled magnetic recording.




Figure 7 – Example of conventional and shingled recording technologies using two zone domains

The ZONE ACTIVATE command (see 5.11) modifies which zones are active zones within a contiguous set of realms.

The result of processing a ZONE ACTIVATE command is that zones in one zone domain within a realm become active zones at the same time as zones in a different zone domain within the same realm become inactive zones.

At any point in time, all the zones in one realm are active zones in only one of the zone domains associated with that realm.

Figure 8 shows an example of the before and after conditions for a ZONE ACTIVATE command that is processed with the AAORB bit set to one (i.e., in the case where realm boundaries affect the processing of a ZONE ACTIVATE command).




Figure 8 – Example of zone activation with the AAORB bit set to one

4.2.4.7 Realm boundary considerations

If the SRB bit (see 5.7.2) is set to zero for a processed zone domain, then the processing of a ZONE ACTIVATE command (see 5.11) shall not result in any realm boundaries changing in that zone domain. If the SRB bit is set to one for a processed zone domain, then updated realm boundary information is available from the REPORT REALMS command (see 5.6) after a ZONE ACTIVATE command is processed to successful completion.

Changes in realm boundary information may result in the RLM ALIGN bit (see 5.11.3) being set to one as an unmet prerequisite for a subsequent ZONE ACTIVATE command, unless the updated information provided by a REPORT REALMS command is used to prepare that ZONE ACTIVATE command.

The realm starting LBAs and ending LBAs returned by a REPORT REALMS command shall not change unless the device server processes a ZONE ACTIVATE command, a FORMAT UNIT command (see SBC-5), or a FORMAT WITH PRESET command (see SBC-5).

4.3 Zone attributes

4.3.1 Zone attributes summary

The zoned block device zone attributes are summarized in table 9.

Table 9 – Summary of zone attributes

Zone attribute	Description	Reference
Zone Type	Type of zone (e.g., CONVENTIONAL)	4.3.2
Zone Condition	Zone's operational characteristics	4.3.3
WPointer	Combination of the write pointer (see 4.5.3.1) and whether the write pointer is valid	4.3.4
RWP Recommended	Indication of whether the processing of a RESET WRITE POINTER command (see 5.9) for this zone is recommended	4.3.5
Non-Sequential Write Resources Active	Indication of whether zone resources are allocated for non-sequential write operations	4.3.6
Predicted Unrecovered Errors Present	Indication of whether predicted unrecovered errors (see SBC-5) are associated with logical blocks in the zone	4.3.7

4.3.2 Zone Type zone attribute

The Zone Type zone attribute (see table 10) indicates the type of zone.

Table 10 – Zone Type zone attribute

Zone Type	Reference
CONVENTIONAL	4.5.2
SEQUENTIAL WRITE PREFERRED	4.5.3.3
SEQUENTIAL WRITE REQUIRED	4.5.3.4
SEQUENTIAL OR BEFORE REQUIRED	4.5.3.5
GAP	4.5.4

The relationships between Zone Type and other zone attributes are shown in table 11.

Table 11 – Relationships between zone attributes

Zone Type	Other zone attributes			
	Non-Sequential Write Resources Active	RWP Recommended	WPointer	Zone Condition
Write pointer zones (see 4.5.3)				
SEQUENTIAL WRITE PREFERRED	see 4.3.6	see 4.3.5	see 4.5.3.1 and table 15	see 4.3.3
SEQUENTIAL WRITE REQUIRED	false ^a			
SEQUENTIAL OR BEFORE REQUIRED	false ^b			
Other zone types				
CONVENTIONAL	false ^c	false ^c	invalid ^c	NOT WRITE POINTER
GAP	false	false	invalid	
^a Zone resources for non-sequential write operations are not allocated for a zone in which all non-sequential write commands are terminated with an error. ^b For a sequential or before required zone (see 4.5.3.5), the values of zone attributes are required to be compatible with the direct access block device type defined in SBC-5. ^c For a conventional zone (see 4.5.2), the values of zone attributes are required to be compatible with the direct access block device type defined in SBC-5.				

4.3.3 Zone Condition zone attribute

The Zone Condition zone attribute (see table 12) is determined by Zone Type (see 4.3.2) and the zone condition state machine state (see 4.5.3.6).

Table 12 – Zone Condition zone attribute

Zone Type	Zone condition state machine	Reference	Zone Condition
Write pointer zones (see 4.5.3)			
SEQUENTIAL WRITE PREFERRED or SEQUENTIAL WRITE REQUIRED	ZC1:Empty state	4.5.3.6.2	EMPTY
	ZC2:Implicit_Open state	4.5.3.6.3	IMPLICITLY OPENED
	ZC3:Explicit_Open state	4.5.3.6.4	EXPLICITLY OPENED
	ZC4:Closed state	4.5.3.6.5	CLOSED
	ZC5:Full state	4.5.3.6.6	FULL
	ZC6:Read_Only state	4.5.3.6.7	READ ONLY
	ZC7:Offline state	4.5.3.6.8	OFFLINE
	ZC8:Inactive	4.5.3.6.9	INACTIVE
SEQUENTIAL OR BEFORE REQUIRED	ZC1:Empty state	4.5.3.6.2	EMPTY
	ZC2:Implicit_Open state	4.5.3.6.3	IMPLICITLY OPENED
	ZC5:Full state	4.5.3.6.6	FULL
	ZC6:Read_Only state	4.5.3.6.7	READ ONLY
	ZC7:Offline state	4.5.3.6.8	OFFLINE
	ZC8:Inactive	4.5.3.6.9	INACTIVE
Other zone types			
CONVENTIONAL	outside the scope of this standard	4.5.2	NOT WRITE POINTER
GAP	outside the scope of this standard	4.5.4	

4.3.4 WPointer zone attribute

For write pointer zones (see 4.5.3), the WPointer zone attribute is the combination of the write pointer (see 4.5.3.1) and whether the write pointer is valid (see table 15).

If Zone Type is CONVENTIONAL or GAP, then WPointer shall be invalid.

4.3.5 RWP Recommended zone attribute

For write pointer zones (see 4.5.3), if RWP Recommended is:

- true, then the device server has determined (e.g., through the detection by the device server that the zone is using too many resources) that the processing of a RESET WRITE POINTER command (see 5.9) specifying this zone is recommended; or

- b) false, then the device server has no recommendation for or against the processing of a RESET WRITE POINTER command specifying this zone.

RWP Recommended shall be set to false if:

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

If RWP Recommended changes from false to true then the device server should establish a unit attention condition for the SCSI initiator port (see SAM-6) associated with each I_T nexus with additional sense code set to ZONE RESET WRITE POINTER RECOMMENDED (see SPC-6) and the INFORMATION field (see SPC-6) set to the zone ID of the zone for which RWP Recommended was set to true.

4.3.6 Non-Sequential Write Resources Active zone attribute

For sequential write preferred zones (see 4.5.3.3), if Non-Sequential Write Resources Active is:

- a) true, then the device server has zone resources allocated for the processing of non-sequential write operations (see 4.5.3.3.2) in this zone; or
- b) false, then the device server does not have zone resources allocated for the processing of non-sequential write operations in this zone.

If a non-sequential write operation is performed in a zone, then the device server may set Non-Sequential Write Resources Active to true. The device server may set Non-Sequential Resources Active 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 write operations in this zone.

If the write pointer indicates the lowest LBA in the zone and the device server has performed any non-sequential write operations 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.

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

- a) the Zone Type is CONVENTIONAL, GAP, SEQUENTIAL WRITE REQUIRED, or SEQUENTIAL OR BEFORE REQUIRED;
- b) the Zone Condition is EMPTY, READ ONLY, OFFLINE, or INACTIVE; or
- c) the device server has not performed a non-sequential write operation in a zone since the last time the Zone Condition was EMPTY.

4.3.7 Predicted Unrecovered Errors Present zone attribute

For conventional zones (see 4.5.2), if Predicted Unrecovered Errors Present is:

- a) true, then the zone contains one or more logical blocks for which:
 - A) read command accesses result in predicted unrecovered read errors (see 4.15.2.3.2); and
 - B) write command accesses result in predicted unrecovered write errors (see 4.15.2.3.3);
 or
- b) false, then none of the logical blocks in the zone are affected by predicted unrecovered errors.

Predicted Unrecovered Errors Present may be changed to:

- a) true by the processing of a REMOVE ELEMENT AND MODIFY ZONES command (see 5.5); and
- b) false by the processing of a RESTORE ELEMENTS AND REBUILD command (see SBC-5).

If Predicted Unrecovered Errors Present is true, then read commands and write commands may report predicted unrecovered errors in ranges of logical blocks as described in SBC-5.

Predicted Unrecovered Errors Present shall be set to false if the Zone Type is SEQUENTIAL WRITE REQUIRED, SEQUENTIAL WRITE PREFERRED, SEQUENTIAL OR BEFORE REQUIRED, or GAP.