



ISO/IEC 14776-323

Edition 1.0 2017-01

INTERNATIONAL STANDARD



Information technology – Small computer system interface (SCSI) –
Part 323: SCSI Block Commands – 3 (SBC-3)

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC 14776-323:2017



THIS PUBLICATION IS COPYRIGHT PROTECTED
Copyright © 2017 ISO/IEC, Geneva, Switzerland

All rights reserved. Unless otherwise specified, no part of this publication may be reproduced or utilized in any form or by any means, electronic or mechanical, including photocopying and microfilm, without permission in writing from either IEC or IEC's member National Committee in the country of the requester. If you have any questions about ISO/IEC copyright or have an enquiry about obtaining additional rights to this publication, please contact the address below or your local IEC member National Committee for further information.

IEC Central Office
3, rue de Varembe
CH-1211 Geneva 20
Switzerland

Tel.: +41 22 919 02 11
Fax: +41 22 919 03 00
info@iec.ch
www.iec.ch

About the IEC

The International Electrotechnical Commission (IEC) is the leading global organization that prepares and publishes International Standards for all electrical, electronic and related technologies.

About IEC publications

The technical content of IEC publications is kept under constant review by the IEC. Please make sure that you have the latest edition, a corrigenda or an amendment might have been published.

IEC Catalogue - webstore.iec.ch/catalogue

The stand-alone application for consulting the entire bibliographical information on IEC International Standards, Technical Specifications, Technical Reports and other documents. Available for PC, Mac OS, Android Tablets and iPad.

IEC publications search - www.iec.ch/searchpub

The advanced search enables to find IEC publications by a variety of criteria (reference number, text, technical committee,...). It also gives information on projects, replaced and withdrawn publications.

IEC Just Published - webstore.iec.ch/justpublished

Stay up to date on all new IEC publications. Just Published details all new publications released. Available online and also once a month by email.

Electropedia - www.electropedia.org

The world's leading online dictionary of electronic and electrical terms containing 20 000 terms and definitions in English and French, with equivalent terms in 15 additional languages. Also known as the International Electrotechnical Vocabulary (IEV) online.

IEC Glossary - std.iec.ch/glossary

65 000 electrotechnical terminology entries in English and French extracted from the Terms and Definitions clause of IEC publications issued since 2002. Some entries have been collected from earlier publications of IEC TC 37, 77, 86 and CISPR.

IEC Customer Service Centre - webstore.iec.ch/csc

If you wish to give us your feedback on this publication or need further assistance, please contact the Customer Service Centre: csc@iec.ch.

STANDARDSISO.COM : Click to view the PDF of IEC 14776-3:2017

INTERNATIONAL STANDARD



**Information technology – Small computer system interface (SCSI) –
Part 323: SCSI Block Commands – 3 (SBC-3)**

INTERNATIONAL
ELECTROTECHNICAL
COMMISSION

ICS 35.200

ISBN 978-2-8322-3721-2

Warning! Make sure that you obtained this publication from an authorized distributor.

FOREWORD	5
INTRODUCTION	7
General	7
SCSI standards family	7
1 Scope	8
2 Normative references	8
3 Terms, definitions, symbols, abbreviations, keywords, and conventions	9
3.1 Terms and definitions	9
3.2 Symbols	19
3.3 Abbreviations	19
3.4 Keywords	20
3.5 Editorial conventions	21
3.6 Numeric and character conventions	22
3.6.1 Numeric conventions	22
3.6.2 Units of measure	22
3.7 State machine conventions	24
4 Direct access block device type model	25
4.1 Direct access block device type model introduction	25
4.2 Direct access block device type model	26
4.2.1 Direct access block device type model overview	26
4.2.2 Logical block access command types	26
4.2.3 Logical block access operation types	26
4.3 Media examples	26
4.3.1 Media examples overview	26
4.3.2 Rotating media	27
4.3.3 Memory media	27
4.4 Removable media	27
4.5 Logical Blocks	28
4.6 Physical blocks	29
4.7 Logical block provisioning	33
4.7.1 Logical block provisioning overview	33
4.7.2 Fully provisioned logical unit	34
4.7.3 Logical block provisioning management	34
4.7.3.1 Logical block provisioning management overview	34
4.7.3.2 Resource provisioned logical unit	34
4.7.3.3 Thin provisioned logical unit	35
4.7.3.4 Unmapping LBAs	35
4.7.3.4.1 Processing unmap requests	35
4.7.3.4.2 Unmap operations	35
4.7.3.4.3 WRITE SAME command and unmap operations	36
4.7.3.5 Autonomous LBA transitions	37
4.7.3.6 Thin provisioned logical unit resource exhaustion considerations	37
4.7.3.7 Logical block provisioning thresholds	37
4.7.3.7.1 Logical block provisioning thresholds overview	37
4.7.3.7.2 Logical block provisioning armed decreasing thresholds	38
4.7.3.7.3 Logical block provisioning armed increasing thresholds	39
4.7.3.7.4 Logical block provisioning threshold notification	40
4.7.4 LBP (logical block provisioning) state machine	40
4.7.4.1 LBP state machine overview	40
4.7.4.2 LBP state machine for thin provisioned logical units supporting anchored LBAs	41
4.7.4.3 LBP state machine for thin provisioned logical units not supporting anchored LBAs	41
4.7.4.4 LBP state machine for resource provisioned logical units	42
4.7.4.5 Performing read operations with respect to logical block provisioning	42
4.7.4.6 LBP1:Mapped state	44

4.7.4.6.1 LBP1:Mapped state description.....	44
4.7.4.6.2 Transition LBP1:Mapped to LBP2:Deallocated	44
4.7.4.6.3 Transition LBP1:Mapped to LBP3:Anchored	44
4.7.4.7 LBP2:Deallocated state	44
4.7.4.7.1 LBP2:Deallocated state description.....	44
4.7.4.7.2 Transition LBP2:Deallocated to LBP1:Mapped	44
4.7.4.7.3 Transition LBP2:Deallocated to LBP3:Anchored	45
4.7.4.8 LBP3:Anchored state	45
4.7.4.8.1 LBP3:Anchored state description	45
4.7.4.8.2 Transition LBP3:Anchored to LBP1:Mapped	45
4.7.4.8.3 Transition LBP3:Anchored to LBP2:Deallocated	45
4.8 Data de-duplication.....	45
4.9 Ready state	45
4.10 Initialization.....	46
4.11 Sanitize operations	46
4.11.1 Sanitize operations overview	46
4.11.2 Performing a sanitize operation	47
4.11.3 Completing a sanitize operation.....	48
4.12 Write protection	49
4.13 Medium defects	49
4.13.1 Medium defects overview	49
4.13.2 Generation of defect lists	52
4.14 Write and unmap failures.....	53
4.15 Caches	53
4.15.1 Caches overview.....	53
4.15.2 Read caching	53
4.15.3 Write caching	53
4.15.4 Command interactions with caches	54
4.15.5 Write operation and write medium operation interactions with caches	54
4.15.6 Read operation and read medium operation interactions with caches	54
4.15.7 Verify medium operation interactions with caches.....	55
4.15.8 Unmap operation interactions with caches	55
4.15.9 Power loss effects on caches	55
4.16 Implicit head of queue command processing	56
4.17 Reservations.....	56
4.18 Error reporting	58
4.18.1 Error reporting overview.....	58
4.18.2 Processing pseudo unrecovered errors	60
4.18.3 Block commands sense data descriptor	61
4.18.4 User data segment referral sense data descriptor.....	62
4.18.5 Direct-access block device sense data descriptor	64
4.19 Model for XOR commands	65
4.19.1 Model for XOR commands overview	65
4.19.2 SCSI storage array device supervised XOR operations	66
4.19.2.1 SCSI storage array device supervised XOR operations overview	66
4.19.2.2 Update write operation	66
4.19.2.3 Regenerate operation.....	66
4.19.2.4 Rebuild operation	67
4.19.3 Array subsystem considerations.....	67
4.19.3.1 Array subsystem considerations overview	67
4.19.3.2 Access to an inconsistent stripe	67
4.20 Rebuild assist mode	67
4.20.1 Rebuild assist mode overview	67
4.20.2 Enabling rebuild assist mode	68
4.20.3 Using the rebuild assist mode.....	68
4.20.3.1 Using rebuild assist mode overview	68
4.20.3.2 Unpredicted unrecovered read error	68
4.20.3.3 Predicted unrecovered read error	69

STALLPDF.COM: Click to view the full PDF of ISO/IEC 14776-323:2017

4.20.3.4 Unpredicted unrecovered write error.....	69
4.20.3.5 Predicted unrecovered write error.....	69
4.20.4 Disabling the rebuild assist mode.....	70
4.20.5 Testing rebuild assist mode.....	70
4.21 START STOP UNIT and power conditions.....	70
4.21.1 START STOP UNIT and power conditions overview.....	70
4.21.2 Processing of concurrent START STOP UNIT commands.....	70
4.21.3 Managing logical block access commands during a change to the active power condition.....	71
4.21.4 Stopped power condition.....	71
4.21.5 START STOP UNIT and power condition state machine.....	71
4.21.5.1 START STOP UNIT and power condition state machine overview.....	71
4.21.5.2 SSU_PC0:Powered_On state.....	73
4.21.5.2.1 SSU_PC0:Powered_On state description.....	73
4.21.5.2.2 Transition SSU_PC0:Powered_On to SSU_PC4:Active_Wait.....	74
4.21.5.2.3 Transition SSU_PC0:Powered_On to SSU_PC8:Stopped.....	74
4.21.5.3 SSU_PC1:Active state.....	74
4.21.5.3.1 SSU_PC1:Active state description.....	74
4.21.5.3.2 Transition SSU_PC1:Active to SSU_PC5:Wait_Idle.....	74
4.21.5.3.3 Transition SSU_PC1:Active to SSU_PC6:Wait_Standby.....	74
4.21.5.3.4 Transition SSU_PC1:Active to SSU_PC10:Wait_Stopped.....	75
4.21.5.4 SSU_PC2:Idle state.....	75
4.21.5.4.1 SSU_PC2:Idle state description.....	75
4.21.5.4.2 Transition SSU_PC2:Idle to SSU_PC4:Active_Wait.....	75
4.21.5.4.3 Transition SSU_PC2:Idle to SSU_PC5:Wait_Idle.....	75
4.21.5.4.4 Transition SSU_PC2:Idle to SSU_PC6:Wait_Standby.....	76
4.21.5.4.5 Transition SSU_PC2:Idle to SSU_PC7:Idle_Wait.....	76
4.21.5.4.6 Transition SSU_PC2:Idle to SSU_PC10:Wait_Stopped.....	76
4.21.5.5 SSU_PC3:Standby state.....	76
4.21.5.5.1 SSU_PC3:Standby state description.....	76
4.21.5.5.2 Transition SSU_PC3:Standby to SSU_PC4:Active_Wait.....	76
4.21.5.5.3 Transition SSU_PC3:Standby to SSU_PC6:Wait_Standby.....	77
4.21.5.5.4 Transition SSU_PC3:Standby to SSU_PC7:Idle_Wait.....	77
4.21.5.5.5 Transition SSU_PC3:Standby to SSU_PC9:Standby_Wait.....	77
4.21.5.5.6 Transition SSU_PC3:Standby to SSU_PC10:Wait_Stopped.....	78
4.21.5.6 SSU_PC4:Active_Wait state.....	78
4.21.5.6.1 SSU_PC4:Active_Wait state description.....	78
4.21.5.6.2 Transition SSU_PC4:Active_Wait to SSU_PC1:Active.....	79
4.21.5.7 SSU_PC5:Wait_Idle state.....	79
4.21.5.7.1 SSU_PC5:Wait_Idle state description.....	79
4.21.5.7.2 Transition SSU_PC5:Wait_Idle to SSU_PC2:Idle.....	79
4.21.5.8 SSU_PC6:Wait_Standby state.....	79
4.21.5.8.1 SSU_PC6:Wait_Standby state description.....	79
4.21.5.8.2 Transition SSU_PC6:Wait_Standby to SSU_PC3:Standby.....	79
4.21.5.9 SSU_PC7:Idle_Wait state.....	79
4.21.5.9.1 SSU_PC7:Idle_Wait state description.....	79
4.21.5.9.2 Transition SSU_PC7:Idle_Wait to SSU_PC2:Idle.....	80
4.21.5.10 SSU_PC8:Stopped state.....	80
4.21.5.10.1 SSU_PC8:Stopped state description.....	80
4.21.5.10.2 Transition SSU_PC8:Stopped to SSU_PC4:Active_Wait.....	80
4.21.5.10.3 Transition SSU_PC8:Stopped to SSU_PC7:Idle_Wait.....	81
4.21.5.10.4 Transition SSU_PC8:Stopped to SSU_PC9:Standby_Wait.....	81
4.21.5.11 SSU_PC9:Standby_Wait state.....	81
4.21.5.11.1 SSU_PC9:Standby_Wait state description.....	81
4.21.5.11.2 Transition SSU_PC9:Standby_Wait to SSU_PC3:Standby.....	81
4.21.5.12 SSU_PC10:Wait_Stopped state.....	82
4.21.5.12.1 SSU_PC10:Wait_Stopped state description.....	82
4.21.5.12.2 Transition SSU_PC10:Wait_Stopped to SSU_PC8:Stopped.....	82
4.22 Protection information model.....	82

4.22.1 Protection information overview.....	82
4.22.2 Protection types	82
4.22.2.1 Protection types overview	82
4.22.2.2 Type 0 protection.....	83
4.22.2.3 Type 1 protection.....	84
4.22.2.4 Type 2 protection.....	84
4.22.2.5 Type 3 protection.....	85
4.22.3 Protection information format.....	85
4.22.4 Logical block guard	89
4.22.4.1 Logical block guard overview	89
4.22.4.2 CRC generation.....	89
4.22.4.3 CRC checking	90
4.22.4.4 CRC test cases	90
4.22.5 Application of protection information.....	90
4.22.6 Protection information and commands	91
4.23 Grouping function	91
4.24 Background scan operations	91
4.24.1 Background scan overview	91
4.24.2 Background pre-scan operations	92
4.24.2.1 Enabling background pre-scan operations.....	92
4.24.2.2 Suspending and resuming background pre-scan operations.....	92
4.24.2.3 Halting background pre-scan operations.....	93
4.24.3 Background medium scan	93
4.24.3.1 Enabling background medium scan operations	93
4.24.3.2 Suspending and resuming background medium scan operations.....	94
4.24.3.3 Halting background medium scan operations	94
4.24.4 Interpreting the logged background scan results.....	95
4.25 Association between commands and CbCS permission bits	95
4.26 Deferred microcode activation.....	97
4.27 Model for uninterrupted sequences on LBA ranges	97
4.28 Referrals	97
4.28.1 Referrals overview	97
4.28.2 Discovering referrals	98
4.28.3 Referrals in sense data	99
4.29 ORWRITE commands	100
4.29.1 ORWRITE commands overview	100
4.29.2 ORWgeneration code.....	100
4.29.2.1 ORWgeneration code overview.....	100
4.29.2.2 ORWgeneration code processing	101
4.29.3 Change generation and clear operation.....	101
4.29.4 Set operation.....	103
4.30 Block device ROD token operations.....	104
4.30.1 Block device ROD token operations overview	104
4.30.2 POPULATE TOKEN command and WRITE USING TOKEN command completion	105
4.30.3 Block device specific ROD tokens	105
4.30.4 Block device zero ROD token	106
4.30.5 ROD token device type specific data	106
5 Commands for direct access block devices	108
5.1 Commands for direct access block devices overview	108
5.2 COMPARE AND WRITE command	111
5.3 FORMAT UNIT command	113
5.3.1 FORMAT UNIT command overview	113
5.3.2 FORMAT UNIT parameter list.....	117
5.3.2.1 FORMAT UNIT parameter list overview.....	117
5.3.2.2 Parameter list header	117
5.3.2.3 Initialization pattern descriptor.....	122
5.4 GET LBA STATUS command	123

- 5.4.1 GET LBA STATUS command overview..... 123
- 5.4.2 GET LBA STATUS parameter data 125
 - 5.4.2.1 GET LBA STATUS parameter data overview 125
 - 5.4.2.2 LBA status descriptor 126
 - 5.4.2.3 LBA status descriptor relationships 126
- 5.5 ORWRITE (16) command 127
- 5.6 ORWRITE (32) command 133
- 5.7 POPULATE TOKEN command 135
 - 5.7.1 POPULATE TOKEN command overview 135
 - 5.7.2 POPULATE TOKEN parameter list..... 136
 - 5.7.3 Block device range descriptor..... 138
- 5.8 PRE-FETCH (10) command..... 139
- 5.9 PRE-FETCH (16) command..... 140
- 5.10 PREVENT ALLOW MEDIUM REMOVAL command 141
- 5.11 READ (10) command 142
- 5.12 READ (12) command 146
- 5.13 READ (16) command 148
- 5.14 READ (32) command 149
- 5.15 READ CAPACITY (10) command 150
 - 5.15.1 READ CAPACITY (10) overview 150
 - 5.15.2 READ CAPACITY (10) parameter data 151
- 5.16 READ CAPACITY (16) command 151
 - 5.16.1 READ CAPACITY (16) command overview..... 151
 - 5.16.2 READ CAPACITY (16) parameter data 152
- 5.17 READ DEFECT DATA (10) command 154
 - 5.17.1 READ DEFECT DATA (10) command overview..... 154
 - 5.17.2 READ DEFECT DATA (10) parameter data 156
- 5.18 READ DEFECT DATA (12) command 156
 - 5.18.1 READ DEFECT DATA (12) command overview..... 156
 - 5.18.2 READ DEFECT DATA (12) parameter data 158
- 5.19 READ LONG (10) command 159
- 5.20 READ LONG (16) command 161
- 5.21 REASSIGN BLOCKS command..... 161
 - 5.21.1 REASSIGN BLOCKS command overview..... 161
 - 5.21.2 REASSIGN BLOCKS parameter list..... 163
- 5.22 RECEIVE ROD TOKEN INFORMATION 165
 - 5.22.1 RECEIVE ROD TOKEN INFORMATION overview 165
 - 5.22.2 RECEIVE ROD TOKEN INFORMATION parameter data for POPULATE TOKEN command..... 165
 - 5.22.3 The RECEIVE ROD TOKEN INFORMATION parameter data for the WRITE USING
TOKEN command..... 168
- 5.23 REPORT REFERRALS command 169
 - 5.23.1 REPORT REFERRALS command overview 169
 - 5.23.2 REPORT REFERRALS parameter data 170
- 5.24 SANITIZE command..... 171
 - 5.24.1 SANITIZE command overview..... 171
 - 5.24.2 SANITIZE command service actions 172
 - 5.24.2.1 SANITIZE command service actions overview 172
 - 5.24.2.2 OVERWRITE service action..... 172
 - 5.24.2.3 BLOCK ERASE service action 173
 - 5.24.2.4 CRYPTOGRAPHIC ERASE service action 174
 - 5.24.2.5 EXIT FAILURE MODE service action 174
- 5.25 START STOP UNIT command..... 174
- 5.26 SYNCHRONIZE CACHE (10) command..... 178
- 5.27 SYNCHRONIZE CACHE (16) command..... 179
- 5.28 UNMAP command..... 180
 - 5.28.1 UNMAP command overview 180
 - 5.28.2 UNMAP parameter list 181
- 5.29 VERIFY (10) command 182

5.30 VERIFY (12) command	195
5.31 VERIFY (16) command	196
5.32 VERIFY (32) command	197
5.33 WRITE (10) command	198
5.34 WRITE (12) command	201
5.35 WRITE (16) command	202
5.36 WRITE (32) command	203
5.37 WRITE AND VERIFY (10) command	204
5.38 WRITE AND VERIFY (12) command	205
5.39 WRITE AND VERIFY (16) command	206
5.40 WRITE AND VERIFY (32) command	207
5.41 WRITE LONG (10) command	208
5.42 WRITE LONG (16) command	211
5.43 WRITE SAME (10) command	212
5.44 WRITE SAME (16) command	214
5.45 WRITE SAME (32) command	215
5.46 WRITE USING TOKEN command	217
5.46.1 WRITE USING TOKEN command overview	217
5.46.2 WRITE USING TOKEN parameter list	218
5.47 XDWRITEREAD (10) command	220
5.48 XDWRITEREAD (32) command	222
5.49 XPWRITE (10) command	222
5.50 XPWRITE (32) command	224
6 Parameters for direct access block devices	225
6.1 Parameters for direct access block devices introduction	225
6.2 Address descriptors	225
6.2.1 Address descriptor overview	225
6.2.2 Short block format address descriptor	226
6.2.3 Extended bytes from index address descriptor	226
6.2.4 Extended physical sector format address descriptor	228
6.2.5 Long block format address descriptor	229
6.2.6 Bytes from index format address descriptor	229
6.2.7 Physical sector format address descriptor	230
6.3 Diagnostic parameters	231
6.3.1 Diagnostic parameters overview	231
6.3.2 Rebuild Assist Input diagnostic page	232
6.3.3 Rebuild Assist Output diagnostic page	233
6.3.4 Translate Address Input diagnostic page	234
6.3.5 Translate Address Output diagnostic page	236
6.4 Log parameters	237
6.4.1 Log parameters overview	237
6.4.1.1 Summary of log pages	237
6.4.1.2 Setting and resetting log parameters	237
6.4.2 Background Scan log page	238
6.4.2.1 Background Scan log page overview	238
6.4.2.2 Background Scan Status log parameter	240
6.4.2.3 Background Scan Results log parameter	242
6.4.3 Format Status log page	245
6.4.3.1 Format Status log page overview	245
6.4.3.2 Format Data Out log parameter	246
6.4.3.3 Grown Defects During Certification log parameter	247
6.4.3.4 Total Blocks Reassigned During Format log parameter	248
6.4.3.5 Total New Blocks Reassigned log parameter	249
6.4.3.6 Power On Minutes Since Format log parameter	250
6.4.4 Logical Block Provisioning log page	251
6.4.4.1 Logical Block Provisioning log page overview	251
6.4.4.2 Available LBA Mapping Resource Count log parameter	253

6.4.4.2.1 Available LBA Mapping Resource Count log parameter overview	253
6.4.4.2.2 RESOURCE COUNT field.....	254
6.4.4.3 Used LBA Mapping Resource Count log parameter	254
6.4.4.4 De-duplicated LBA Resource Count log parameter	255
6.4.4.5 Compressed LBA Resource Count log parameter	256
6.4.4.6 Total Efficiency LBA Resource Count log parameter	257
6.4.5 Non-volatile Cache log page.....	258
6.4.5.1 Non-volatile Cache log page overview	258
6.4.5.2 Remaining Nonvolatile Time log parameter	259
6.4.5.3 Maximum Nonvolatile Time log parameter	260
6.4.6 Solid State Media log page	260
6.4.6.1 Solid State Media log page overview	260
6.4.6.2 Percentage Used Endurance Indicator log parameter	262
6.5 Mode parameters	263
6.5.1 Mode parameters overview.....	263
6.5.2 Mode parameter block descriptors.....	264
6.5.2.1 Mode parameter block descriptors overview	264
6.5.2.2 Short LBA mode parameter block descriptor	264
6.5.2.3 Long LBA mode parameter block descriptor	266
6.5.3 Application Tag mode page	267
6.5.3.1 Introduction.....	267
6.5.3.2 Application tag descriptor	269
6.5.4 Background Control mode page	270
6.5.5 Caching mode page.....	272
6.5.6 Informational Exceptions Control mode page	276
6.5.7 Logical Block Provisioning mode page	281
6.5.7.1 Logical Block Provisioning mode page overview	281
6.5.7.2 Threshold descriptor format	282
6.5.8 Read-Write Error Recovery mode page.....	283
6.5.9 Verify Error Recovery mode page.....	289
6.6 Vital product data (VPD) parameters.....	290
6.6.1 VPD parameters overview	290
6.6.2 Block Device Characteristics VPD page	291
6.6.3 Block Limits VPD page	294
6.6.4 Logical Block Provisioning VPD page.....	297
6.6.5 Referrals VPD page	299
6.6.6 Third-Party Copy VPD page	300
6.6.6.1 Third-Party Copy VPD page overview.....	300
6.6.6.2 Block device third-party copy descriptor type codes	300
6.6.6.3 Block Device ROD Token Limits descriptor	301
6.7 Copy manager parameters.....	302
Annex A (informative) Numeric order codes	303
A.1 Variable length CDBs	303
A.2 Service action CDBs	304
Annex B (informative) XOR command examples.....	305
B.1 XOR command examples overview	305
B.2 Update write operation	305
B.3 Regenerate operation	306
B.4 Rebuild operation.....	307
Annex C (informative) CRC example in C.....	309
Annex D (informative) Sense information for locked or encrypted logical units	311
Annex E (informative) Optimizing block access characteristics	312
E.1 Optimizing block access overview	312

E.2 Starting logical block offset	312
E.3 Optimal granularity sizes	312
E.4 Optimizing transfers	312
E.5 Examples	313
Annex F (informative) Logical block provisioning reporting examples	314
F.1 Logical block provisioning reporting examples overview	314
F.2 Interpreting log parameter counts	314
F.3 Dedicated resource, threshold set tracked example	315
F.3.1 Dedicated resource, threshold set tracked example overview	315
F.3.2 Dedicated resource, threshold set tracked example configuration	315
F.3.3 Dedicated resource, threshold set tracked example sequence	316
F.3.4 Dedicated resource, threshold set tracked example initial conditions	317
F.3.5 Operations that occur	317
F.3.6 Dedicated resource, threshold set tracked example final log page values	318
F.4 Shared resource, logical block tracked example	318
F.4.1 Shared resource, logical block tracked example overview	318
F.4.2 Shared resource, logical block tracked example configuration	319
F.4.3 Shared resource, logical block tracked example time line	319
F.4.4 Shared resource, logical block tracked example initial conditions	320
F.4.5 Operations that occur	320
F.4.6 Shared resource, logical block tracked example final log page values	321
F.5 Shared available, dedicated used, logical block tracked example	322
F.5.1 Shared available, dedicated used, logical block tracked example overview	322
F.5.2 Shared available, dedicated used, logical block tracked example configuration	322
F.5.3 Shared available, dedicated used, logical block tracked example time line	322
F.5.4 Shared available, dedicated used, logical block tracked example initial conditions	323
F.5.5 Operations that occur	323
F.5.6 Shared available, dedicated used, example final log page values	324
Annex G (informative) Discovering referrals examples	325
G.1 Referrals example with no user data segment multiplier	325
G.2 Referrals example with non-zero user data segment multiplier	327
Bibliography	329

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC 14776-323:2017

Figure 0 – SCSI document relationships	7
Figure 1 – Example state machine figure	24
Figure 2 – One or more physical blocks per logical block examples	30
Figure 3 – One or more logical blocks per physical block examples	31
Figure 4 – Two logical blocks per physical block alignment examples	31
Figure 5 – Four logical blocks per physical block alignment examples	32
Figure 6 – Examples of the relationship between mapped and unmapped LBAs and physical blocks	33
Figure 7 – Armed decreasing threshold operation	39
Figure 8 – Armed increasing threshold operation	39
Figure 9 – LBP state machine (anchored LBAs supported and deallocated LBAs supported)	41
Figure 10 – LBP state machine (anchored LBAs not supported)	42
Figure 11 – LBP state machine (deallocated LBAs not supported)	42
Figure 12 – SSU_PC state machine	73
Figure 13 – Referrals	98
Figure B.1 – Update write operation (SCSI storage array device supervised)	306
Figure B.2 – Regenerate operation (SCSI storage array device supervised)	307
Figure B.3 – Rebuild operation (SCSI storage array device supervised)	308
Figure G.1 – Referrals example with no user data segment multiplier	325
Figure G.2 – Referrals example with non-zero user data segment multiplier	327

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC 14776-323:2017

Table 1 – Numbering convention examples	22
Table 2 – Comparison of decimal prefixes and binary prefixes	23
Table 3 – Direct access block device type model topics	25
Table 4 – Logical block provisioning states supported by logical block provisioning type	33
Table 5 – WRITE SAME command and unmap operations	36
Table 6 – Threshold resource value, threshold type value, and threshold arming value for logical block provisioning thresholds	38
Table 7 – Logical block data returned by a read operation from a mapped LBA	43
Table 8 – Logical block data returned by a read operation from an unmapped LBA	43
Table 9 – Defect lists (i.e., PLIST and GLIST)	50
Table 10 – Address descriptor formats	52
Table 11 – SBC-3 commands that are allowed in the presence of various reservations	57
Table 12 – Example error conditions	59
Table 13 – Sense data field usage for direct access block devices	60
Table 14 – Block commands sense data descriptor format	62
Table 15 – User data segment referral sense data descriptor format	62
Table 16 – User data segment referral descriptor format	63
Table 17 – Target port group descriptor	64
Table 18 – Direct-access block device sense data descriptor format	65
Table 19 – Summary of states in the SSU_PC state machine	72
Table 20 – Logical block data format with a single protection information interval	85
Table 21 – An example of the logical block data for a logical block with more than one protection information interval	86
Table 22 – Content of the first LOGICAL BLOCK REFERENCE TAG field for the first logical block in the Data-In Buffer and/or Data-Out Buffer	87
Table 23 – Content of subsequent LOGICAL BLOCK REFERENCE TAG fields for a logical block in the Data-In Buffer and/or Data-Out Buffer	88
Table 24 – CRC polynomials	89
Table 25 – CRC test cases	90
Table 26 – Associations between commands and CbCS permissions	96
Table 27 – Commands that require uninterrupted sequences	97
Table 28 – Performing an ORWRITE set operation	103
Table 29 – ROD token type values	106
Table 30 – Block device zero ROD token format	106
Table 31 – Commands for direct access block devices	108
Table 32 – COMPARE AND WRITE command	112
Table 33 – FORMAT UNIT command	114
Table 34 – FORMAT UNIT command address descriptor support requirements	116
Table 35 – FORMAT UNIT parameter list	117
Table 36 – Short parameter list header	117
Table 37 – Long parameter list header	118
Table 38 – FMTPINFO field and PROTECTION FIELD USAGE field	119
Table 39 – Initialization pattern descriptor	122
Table 40 – INITIALIZATION PATTERN TYPE field	123
Table 41 – GET LBA STATUS command	124
Table 42 – GET LBA STATUS parameter data	125
Table 43 – LBA status descriptor format	126
Table 44 – PROVISIONING STATUS field	126
Table 45 – ORWRITE (16) command	127
Table 46 – ORPROTECT field - checking protection information from the read operations	128
Table 47 – ORPROTECT field - checking protection information from the Data-Out Buffer	131
Table 48 – ORWRITE (32) command	133
Table 49 – BMOP field	134
Table 50 – POPULATE TOKEN command	135
Table 51 – POPULATE TOKEN parameter list	136
Table 52 – Block device range descriptor	138
Table 53 – PRE-FETCH (10) command	139
Table 54 – PRE-FETCH (16) command	140

Table 55 – PREVENT ALLOW MEDIUM REMOVAL command	141
Table 56 – PREVENT field	141
Table 57 – READ (10) command	142
Table 58 – RDPROTECT field	143
Table 59 – READ (12) command	147
Table 60 – READ (16) command	148
Table 61 – READ (32) command	149
Table 62 – READ CAPACITY (10) command	150
Table 63 – READ CAPACITY (10) parameter data	151
Table 64 – READ CAPACITY (16) command	152
Table 65 – READ CAPACITY (16) parameter data	152
Table 66 – P_TYPE field and PROT_EN bit	153
Table 67 – LOGICAL BLOCKS PER PHYSICAL BLOCK EXPONENT field	153
Table 68 – READ DEFECT DATA (10) command	154
Table 69 – REQ_PLIST bit and REQ_GLIST bit	155
Table 70 – READ DEFECT DATA (10) parameter data	156
Table 71 – READ DEFECT DATA (12) command	157
Table 72 – READ DEFECT DATA (12) parameter data	158
Table 73 – READ LONG (10) command	159
Table 74 – READ LONG (16) command	161
Table 75 – REASSIGN BLOCKS command	162
Table 76 – REASSIGN BLOCKS parameter list	163
Table 77 – REASSIGN BLOCKS short parameter list header	163
Table 78 – REASSIGN BLOCKS long parameter list header	163
Table 79 – Reassign LBA if the LONGLBA bit is set to zero	164
Table 80 – Reassign LBA if the LONGLBA bit is set to one	164
Table 81 – RECEIVE ROD TOKEN INFORMATION reference	165
Table 82 – RECEIVE ROD TOKEN INFORMATION parameter data for POPULATE TOKEN	166
Table 83 – RECEIVE ROD TOKEN INFORMATION parameter data for WRITE USING TOKEN	168
Table 84 – REPORT REFERRALS command	169
Table 85 – REPORT REFERRALS parameter data	170
Table 86 – SANITIZE command	171
Table 87 – SANITIZE service action codes	172
Table 88 – OVERWRITE service action parameter list	172
Table 89 – TEST field	173
Table 90 – START STOP UNIT command	175
Table 91 – POWER CONDITION and POWER CONDITION MODIFIER field	176
Table 92 – SYNCHRONIZE CACHE (10) command	178
Table 93 – SYNCHRONIZE CACHE (16) command	179
Table 94 – UNMAP command	180
Table 95 – UNMAP parameter list	181
Table 96 – UNMAP block descriptor	182
Table 97 – Data-Out Buffer contents for the VERIFY (10) command	183
Table 98 – VERIFY (10) command	183
Table 99 – VRPROTECT field with the BYTCHK field set to 00b – checking protection information from the verify operations	185
Table 100 – VRPROTECT field with the BYTCHK field set to 01b or 11b – checking protection information from the verify operations	188
Table 101 – VRPROTECT field with the BYTCHK field set to 01b or 11b – checking protection information from the Data-Out Buffer	190
Table 102 – VRPROTECT field with the BYTCHK field set to 01b or 11b – compare operation requirements ..	192
Table 103 – VERIFY (12) command	195
Table 104 – VERIFY (16) command	196
Table 105 – VERIFY (32) command	197
Table 106 – WRITE (10) command	198
Table 107 – WRPROTECT field	199
Table 108 – WRITE (12) command	201
Table 109 – WRITE (16) command	202

Table 110 – WRITE (32) command	203
Table 111 – WRITE AND VERIFY (10) command	204
Table 112 – WRITE AND VERIFY (12) command	205
Table 113 – WRITE AND VERIFY (16) command	206
Table 114 – WRITE AND VERIFY (32) command	207
Table 115 – WRITE LONG (10) command	208
Table 116 – COR_DIS bit, WR_UNCOR bit, and PBLOCK bit	209
Table 117 – WRITE LONG (16) command	211
Table 118 – WRITE SAME (10) command	213
Table 119 – UNMAP bit, ANCHOR bit, and ANC_SUP bit relationships	214
Table 120 – WRITE SAME (16) command	215
Table 121 – NDOB bit and UNMAP bit interactions	215
Table 122 – WRITE SAME (32) command	216
Table 123 – WRITE USING TOKEN command	217
Table 124 – WRITE USING TOKEN parameter list	218
Table 125 – XDWRITEREAD (10) command	221
Table 126 – XDWRITEREAD (32) command	222
Table 127 – XPWRITE (10) command	223
Table 128 – XPWRITE (32) command	224
Table 129 – Parameters for direct access block devices	225
Table 130 – Address descriptors	226
Table 131 – Short block format address descriptor (000b)	226
Table 132 – Extended bytes from index format address descriptor (001b)	227
Table 133 – Sorting order for extended bytes from index format address descriptors	227
Table 134 – Extended physical sector format address descriptor (010b)	228
Table 135 – Sorting order for extended physical sector format address descriptors	229
Table 136 – Long block format address descriptor (011b)	229
Table 137 – Bytes from index format address descriptor (100b)	229
Table 138 – Sorting order for bytes from index format address descriptors	230
Table 139 – Physical sector format address descriptor (101b)	230
Table 140 – Sorting order for physical sector format address descriptors	230
Table 141 – Diagnostic page codes for direct access block devices	231
Table 142 – Rebuild Assist Input diagnostic page	232
Table 143 – Rebuild Assist Output diagnostic page	233
Table 144 – Translate Address Input diagnostic page	234
Table 145 – Translate Address Output diagnostic page	236
Table 146 – Log page codes and subpage codes for direct access block devices	237
Table 147 – Keywords for resetting or changing log parameters	238
Table 148 – Background Scan log page parameter codes	238
Table 149 – Background Scan log page	239
Table 150 – Background Scan Status log parameter format	240
Table 151 – BACKGROUND SCAN STATUS field	241
Table 152 – Background Scan Results log parameter format	242
Table 153 – REASSIGN STATUS field	243
Table 154 – Format Status log page parameter codes	245
Table 155 – Format Status log page	245
Table 156 – Format Data Out log parameter format	246
Table 157 – Grown Defects During Certification log parameter format	247
Table 158 – Total Blocks Reassigned During Format log parameter format	248
Table 159 – Total New Blocks Reassigned log parameter format	249
Table 160 – Power On Minutes Since Format log parameter format	250
Table 161 – Logical Block Provisioning log parameters	251
Table 162 – Logical Block Provisioning log page	252
Table 163 – Available LBA Mapping Resource Count log parameter format	253
Table 164 – SCOPE field	253
Table 165 – Used LBA Mapping Resource Count log parameter format	254
Table 166 – De-duplicated LBA Resource Count log parameter format	255
Table 167 – Compressed LBA Resource Count log parameter format	256

Table 168 – Total Efficiency LBA Resource Count log parameter format	257
Table 169 – Nonvolatile Cache log parameters	258
Table 170 – Nonvolatile Cache log page	258
Table 171 – Remaining Nonvolatile Time parameter data	259
Table 172 – REMAINING NONVOLATILE TIME field	259
Table 173 – Maximum Nonvolatile Time parameter data	260
Table 174 – MAXIMUM NONVOLATILE TIME field	260
Table 175 – Solid State Media log parameters	261
Table 176 – Solid State Media log page	261
Table 177 – Percentage Used Endurance Indicator log parameter format	262
Table 178 – Mode page codes and subpage codes for direct access block devices	263
Table 179 – DEVICE-SPECIFIC PARAMETER field for direct access block devices	264
Table 180 – Short LBA mode parameter block descriptor	265
Table 181 – Long LBA mode parameter block descriptor	266
Table 182 – Application Tag mode page	268
Table 183 – Application tag descriptor format	269
Table 184 – Background Control mode page	270
Table 185 – Caching mode page	272
Table 186 – DEMAND READ RETENTION PRIORITY field	273
Table 187 – WRITE RETENTION PRIORITY field	274
Table 188 – SYNC_PROG field	275
Table 189 – Informational Exceptions Control mode page	276
Table 190 – Definitions for the combinations of values in EWASC, DEXCPT, and TEST	277
Table 191 – Method of reporting informational exceptions (MRIE) field	278
Table 192 – Use of the INTERVAL TIMER field and the REPORT COUNT field based on the MRIE field	280
Table 193 – Logical Block Provisioning mode page	281
Table 194 – Threshold descriptor format	282
Table 195 – THRESHOLD TYPE field	282
Table 196 – THRESHOLD ARMING field	282
Table 197 – Read-Write Error Recovery mode page	283
Table 198 – Error recovery bit combinations	285
Table 199 – Verify Error Recovery mode page	289
Table 200 – VPD page codes for direct access block devices	290
Table 201 – Block Device Characteristics VPD page	291
Table 202 – MEDIUM ROTATION RATE field	291
Table 203 – PRODUCT TYPE field	292
Table 204 – WABEREQ field	292
Table 205 – WACEREQ field	293
Table 206 – NOMINAL FORM FACTOR field	293
Table 207 – Block Limits VPD page	294
Table 208 – Transfer limits for commands	295
Table 209 – Logical Block Provisioning VPD page	297
Table 210 – PROVISIONING TYPE field	298
Table 211 – Referrals VPD page	299
Table 212 – Block device third-party copy descriptor type codes	300
Table 213 – Block Device ROD Token Limits descriptor	301
Table 214 – ROD token device type specific data	302
Table A.1 – Variable length command service action code assignments	303
Table A.2 – SERVICE ACTION IN (16) service actions	304
Table A.3 – SERVICE ACTION OUT (16) service actions	304
Table D.1 – Sense information for locked or encrypted logical units	311
Table F.1 – Dedicated resource, threshold set tracked example capacity information	315
Table F.2 – Dedicated resource, threshold set tracked example capacity information	316
Table F.3 – Dedicated resource, threshold set tracked example initial conditions	317
Table F.4 – Dedicated resource, threshold set tracked example final log page values	318
Table F.5 – Shared resource, logical block tracked example capacity information	319
Table F.6 – Shared resource, logical block tracked example initial conditions	320
Table F.7 – Shared resource, logical block tracked example final log page values	321

Table F.8 – Shared available, dedicated used example capacity information	322
Table F.9 – Shared resource, logical block tracked example initial conditions	323
Table F.10 – Shared available, dedicated used example final log page values	324
Table G.1 – Referrals application client information with no user data segment multiplier	326
Table G.2 – User data segment calculations with no user data segment multiplier	326
Table G.3 – Referrals application client information with non-zero user data segment multiplier	328
Table G.4 – User data segment calculations with non-zero user data segment multiplier	328

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC 14776-323:2017

INFORMATION TECHNOLOGY – SMALL COMPUTER SYSTEM INTERFACE (SCSI) -

Part 323: SCSI Block Commands – 3 (SBC-3)

FOREWORD

- 1) 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. In the field of information technology, ISO and IEC have established a joint technical committee, ISO/IEC JTC 1.
- 2) The formal decisions or agreements of IEC and ISO on technical matters express, as nearly as possible, an international consensus of opinion on the relevant subjects since each technical committee has representation from all interested IEC National Committees and ISO member bodies.
- 3) IEC, ISO and ISO/IEC publications have the form of recommendations for international use and are accepted by IEC National Committees and ISO member bodies in that sense. While all reasonable efforts are made to ensure that the technical content of IEC, ISO and ISO/IEC publications is accurate, IEC or ISO cannot be held responsible for the way in which they are used or for any misinterpretation by any end user.
- 4) In order to promote international uniformity, IEC National Committees and ISO member bodies undertake to apply IEC, ISO and ISO/IEC publications transparently to the maximum extent possible in their national and regional publications. Any divergence between any ISO, IEC or ISO/IEC publication and the corresponding national or regional publication should be clearly indicated in the latter.
- 5) ISO and IEC do not provide any attestation of conformity. Independent certification bodies provide conformity assessment services and, in some areas, access to IEC marks of conformity. ISO or IEC are not responsible for any services carried out by independent certification bodies.
- 6) All users should ensure that they have the latest edition of this publication.
- 7) No liability shall attach to IEC or ISO or its directors, employees, servants or agents including individual experts and members of their technical committees and IEC National Committees or ISO member bodies for any personal injury, property damage or other damage of any nature whatsoever, whether direct or indirect, or for costs (including legal fees) and expenses arising out of the publication of, use of, or reliance upon, this ISO/IEC publication or any other IEC, ISO or ISO/IEC publications.
- 8) Attention is drawn to the normative references cited in this publication. Use of the referenced publications is indispensable for the correct application of this publication.
- 9) Attention is drawn to the possibility that some of the elements of this ISO/IEC publication may be the subject of patent rights. ISO and IEC shall not be held responsible for identifying any or all such patent rights.

International Standard ISO/IEC 14776-323 was prepared by subcommittee 25: Interconnection of information technology equipment, of ISO/IEC joint technical committee 1: Information technology.

The list of all currently available parts of the ISO/IEC 14776 series, under the general title *Information technology – Small computer system interface (SCSI)*, can be found on the IEC web site.

This International Standard has been approved by vote of the member bodies and the voting results may be obtained from the address given on the second title page.

This publication has been drafted in accordance with the ISO/IEC Directives, Part 2, except as described in 3.5 and 3.6.

IMPORTANT - The 'colour inside' logo on the cover page of this publication indicates that it contains colours which are considered to be useful for the correct understanding of its contents. Users should therefore print this document using a colour printer.

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC 14776-323:2017

INTRODUCTION

General

The 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 - 4 (SPC-4) to facilitate operation of SCSI direct-access block devices (e.g., hard disk drives).

SCSI standards family

Figure 0 shows the relationship of this standard to the other standards and related projects in the SCSI family of standards as of the publication of this standard.

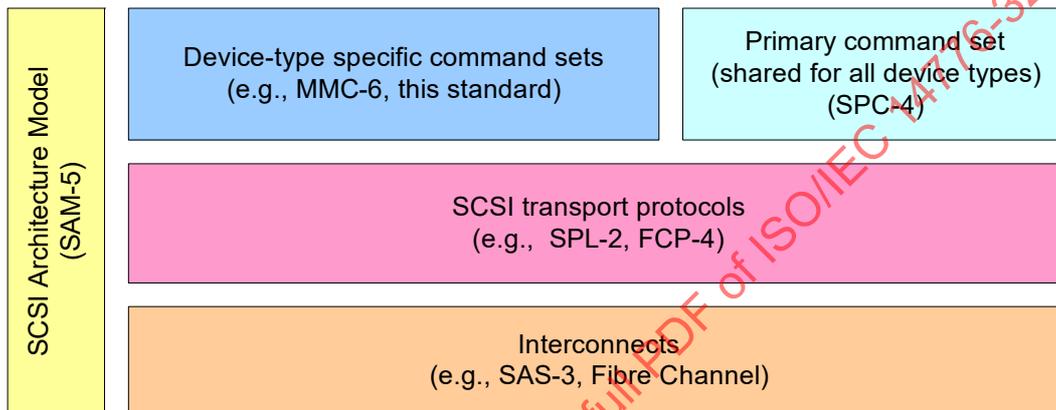


Figure 0 – SCSI document relationships

Figure 0 gives the general relationship of the documents to one another and is not intended to imply a relationship such as a hierarchy, protocol stack, or system architecture.

The set of SCSI standards specifies the interfaces, functions, and operations necessary to ensure interoperability between conforming SCSI implementations. This standard is a functional description. Conforming implementations may employ any design technique that does not violate interoperability. See SAM-5 for more information about the relationships between the SCSI standards.

This standard makes obsolete the following concepts from SBC-2:

- a) linked commands;
- b) the partial medium indicator (PMI) bit and the LOGICAL BLOCK ADDRESS field in the READ CAPACITY (10) command and the READ CAPACITY (16) command;
- c) the READ (6) command and the WRITE (6) command;
- d) the XDREAD (10) command, the XDREAD (32) command, the XDWRITE (10) command, and the XDWRITE (32) command;
- e) the SYNC_NV bit in the SYNCHRONIZE CACHE commands;
- f) the FUA_NV bit in read commands;
- g) the FUA_NV bit in write commands;
- h) the LBDATA bit and the PBDATA bit in the WRITE SAME commands;
- i) the initialization pattern modifier (IP MODIFIER) field in the initialization pattern descriptor in the FORMAT UNIT command; and
- j) the XOR Control mode page.

INFORMATION TECHNOLOGY – SMALL COMPUTER SYSTEM INTERFACE (SCSI) –

Part 323: SCSI Block Commands - 3 (SBC-3)

1 Scope

This part of ISO/IEC 14776 defines the command set extensions to facilitate operation of SCSI direct access block devices. The clauses in this standard, implemented in conjunction with the applicable clauses of SPC-4, specify the standard command set for SCSI direct access block devices.

The objectives of this standard are to:

- a) permit an application client to communicate over a SCSI service delivery subsystem (see SAM-5) with a logical unit that declares itself to be a direct access block device in the PERIPHERAL DEVICE TYPE field of the standard INQUIRY data (see SPC-4); and
- b) define commands and parameters unique to the direct access block device type.

2 Normative references

The following documents are referred to in the text in such a way that some or all of their content constitutes requirements of this document. For dated references, only the edition cited applies. For undated references, the latest edition of the referenced document (including any amendments) applies.

ISO/IEC 14776-262, *Information Technology - Small Computer System Interface (SCSI) - Part 262: SAS Protocol Layer - 2 (SPL-2)*

ISO/IEC 14776-342, *Information Technology - Small Computer System Interface (SCSI) - Part 342: Controller Commands - 2 (SCC-2)*

INCITS 513-2015, *Information Technology - SCSI Primary Commands - 4 (SPC-4)*

INCITS 515-2016, *Information Technology - SCSI Architecture Model - 5 (SAM-5)*

INCITS 517-2015, *Information Technology - SCSI / ATA Translation - 3 (SAT-3)*

INCITS 448-2008, *Information Technology - SCSI Enclosure Services - 2 (SES-2)*

3 Terms, definitions, symbols, abbreviations, keywords, and conventions

3.1 Terms and definitions

For the purposes of this document the following terms and definitions apply.

ISO and IEC maintain terminological databases for use in standardization at the following addresses:

- IEC Electropedia: available at <http://www.electropedia.org/>
- ISO Online browsing platform: available at <http://www.iso.org/obp>

3.1.1

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-4.

3.1.2

AND

Boolean arithmetic function (see 3.1.10) on two binary input values that results in an output value of one if both of the input values are one or zero if either of the input values is zero

3.1.3

AND operation

performance of an AND (see 3.1.2) bitwise on two multiple-bit input values both having the same number of bits

Note 1 to entry: An example of multiple-bit inputs is the current content of a logical block and the content contained in the Data-Out Buffer having the same number of bytes.

3.1.4

anchored

logical block provisioning state of an LBA (see 4.7.1) in which physical capacity has been reserved for the referenced logical block (see 4.7.4.8)

3.1.5

application client

object that is the source of SCSI commands

Note 1 to entry: See SAM-5.

3.1.6

automatic read reassignment

sequence after the device server detects a recovered read error during which the device server, without intervention from an application client, performs a reassign operation on the LBA for which the error was detected

3.1.7

automatic write reassignment

sequence after the device server detects a recovered error or an unrecovered error during which the device server, without intervention from an application client, performs a reassign operation on the LBA for which the error was detected

3.1.8

background function

either a background scan operation (see 4.24) or a device specific background function (see 3.1.23)

3.1.9

bitmap buffer

temporary buffer within a device server (e.g., for one or more bytes of the result of an AND operation (see 3.1.3) or an OR operation (see 3.1.51))

3.1.10**Boolean arithmetic function**

function that produces an output from one or more inputs according to the rules of AND (see 3.1.2), XOR (see 3.1.27), and OR (see 3.1.50)

3.1.11**byte**

sequence of eight contiguous bits considered as a unit

3.1.12**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

3.1.13**check data**

information contained within a redundancy group (see 3.1.73) that may allow lost or destroyed XOR-protected data (see 3.1.105) to be recreated

3.1.14**command**

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

Note 1 to entry: See SAM-5.

3.1.15**command descriptor block****CDB**

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

Note 1 to entry: See SPC-4.

3.1.16**compare operation**

process by which a device server compares two sets of data for equality

3.1.17**cyclic redundancy check****CRC**

error checking mechanism that checks data integrity by computing a polynomial algorithm based checksum (see 4.22.4)

3.1.18**Data-In Buffer**

buffer specified by the application client to receive data from the device server during the processing of a command

Note 1 to entry: See SAM-5 and SPC-4.

3.1.19**Data-Out Buffer**

buffer specified by the application client to supply data that is sent from the application client to the device server during the processing of a command

Note 1 to entry: See SAM-5 and SPC-4.

3.1.20**deallocated**

logical block provisioning state of an LBA (see 4.7.1) in which physical capacity has not been reserved for the referenced logical block (see 4.7.4.7)

3.1.21

defect list

GLIST (see 4.13) or PLIST (see 4.13)

3.1.22

device server

object within a logical unit (see 3.1.43) that processes SCSI commands according to the rules of command management

Note 1 to entry: See SAM-5.

3.1.23

device specific background functions

SCSI target device specific functions that a SCSI target device may perform that have no specific association with application client-initiated operations

Note 1 to entry: See SPC-4.

3.1.24

device type

type of peripheral device (i.e., device model) implemented by the device server and indicated to the application client by the PERIPHERAL DEVICE TYPE field of the standard INQUIRY data (see SPC-4)

3.1.25

direct access block device

device that is capable of containing data stored in logical blocks that each have a unique LBA (see 4.2)

3.1.26

error correcting code

ECC

error checking mechanism that checks data integrity and enables some errors in the logical block data to be corrected

3.1.27

exclusive-or

XOR

Boolean arithmetic function on two binary input values that results in an output value of one if one and only one of the input values is one, or zero if both of the input values are either zero or one

3.1.28

extent

set of logical blocks occupying contiguous LBAs on a logical unit

3.1.29

field

group of one or more contiguous bits, a part of a larger structure (e.g., a CDB (see 3.1.15) or sense data (see SPC-4))

3.1.30

format corrupt

vendor specific condition in which the device server may not be able to perform logical block access commands (see 4.10)

3.1.31

format operation

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

Note 1 to entry: See 4.10.

3.1.32**fully provisioned logical unit**

logical unit that stores logical block data for every LBA and has assigned physical capacity for every LBA

Note 1 to entry: See 4.7.2.

3.1.33**grown defect list****GLIST**

list of physical blocks that the device server has detected as containing medium defects or that the application client has specified as containing medium defects

Note 1 to entry: See 4.13.

3.1.34**hard reset**

condition resulting from the events defined by SAM-5 during which the SCSI device performs the hard reset operations described in SAM-5, this standard, and other applicable command standards (see table 31)

3.1.35**I_T nexus**

relationship between a SCSI initiator port and a SCSI target port

Note 1 to entry: See SAM-5.

3.1.36**I_T nexus loss**

condition resulting from the events defined by SAM-5 during which the SCSI device performs the I_T nexus loss operations described in SAM-5, this standard, and other applicable command standards (see table 31)

3.1.37**LBA mapping resource**

resource used by a logical unit that supports logical block provisioning management

Note 1 to entry: An example of a mapping resource is a physical block or a data structure associated with tracking resource usage.

3.1.38**logical block**

set of data bytes accessed and referenced as a unit (see 4.5)

3.1.39**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 4.2.2.

3.1.40**logical block address****LBA**

value used to reference a logical block (see 4.5)

3.1.41**logical block data**

user data and protection information, if any

3.1.42**logical block length**

number of bytes of user data in a logical block (see 4.5)

3.1.43

logical unit

externally addressable entity within a SCSI target device (see 3.1.76) that implements a SCSI device model

Note 1 to entry: See SAM-5.

3.1.44

logical unit reset

condition resulting from the events defined by SAM-5 in which the logical unit performs the logical unit reset operations described in SAM-5, this standard, and other applicable command standards (see table 31)

3.1.45

mapped

logical block provisioning state of an LBA (see 4.7.1) in which physical capacity has been assigned to the referenced logical block (see 4.7.4.6)

3.1.46

medium

material that is not cache on which data is stored

Note 1 to entry: The plural of medium is media.

Note 2 to entry: An example of a medium in which data is stored is a magnetic disk.

3.1.47

medium defect

area of the medium that results in a recovered error or an unrecovered error when a read medium operation or a write medium operation is performed

Note 1 to entry: See 4.13.

3.1.48

non-volatile cache

cache that retains logical block data through any power cycle

3.1.49

non-volatile medium

medium that retains logical block data through any power cycle

3.1.50

OR

Boolean arithmetic function (see 3.1.10) on two binary input values that results in an output value of one if either of the input values are one or zero if both of the input values are zero

3.1.51

OR operation

performance of an OR (see 3.1.50) bitwise on two multiple-bit input values both having the same number of bits

Note 1 to entry: An example of multiple-bit inputs is the current content of a logical block and the content contained in the Data-Out Buffer having the same number of bytes.

3.1.52

point in time ROD token

ROD token with a ROD type that is a point in time copy ROD

Note 1 to entry: See SPC-4.

3.1.53

physical block

set of data bytes accessed as a unit by the device server (see 4.6)

3.1.54**physical block length**

number of bytes of logical block data in a physical block (see 4.6)

3.1.55**physical element**

component that provides non-volatile storage for an associated group of logical blocks (see 4.20)

3.1.56**power cycle**

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

3.1.57**power on**

condition resulting from the events defined by SAM-5 during which a SCSI device performs the power on operations described in SAM-5, this standard, and other applicable command standards (see table 31)

3.1.58**primary defect list****PLIST**

list of physical blocks containing medium defects that are considered permanent

Note 1 to entry: See 4.13.

3.1.59**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 4.22.

3.1.60**protection information interval**

length of user data that occurs within a logical block before each protection information

3.1.61**pseudo read data**

indeterminate logical block data

3.1.62**pseudo unrecovered error**

simulated error (e.g., created by a WRITE LONG command (see 5.41 and 5.42)) for which a device server reports that it is unable to read or write a logical block, regardless of whether the data on the medium is valid, recoverable, or unrecoverable

3.1.63**pseudo unrecovered error with correction enabled**

pseudo unrecovered error (see 3.1.62) for which a device server performs the maximum error recovery as specified in the Read-Write Error Recovery mode page (see 6.5.8)

Note 1 to entry: See 4.18.2.

3.1.64**pseudo unrecovered error with correction disabled**

pseudo unrecovered error (see 3.1.62) for which a device server performs no error recovery

Note 1 to entry: See 4.18.2.

3.1.65

read cache operation

process by which a device server reads logical blocks for one or more LBAs from cache as described in 4.15

3.1.66

read command

command that requests read operations

Note 1 to entry: See 4.2.2.

3.1.67

read medium operation

process by which a device server reads logical blocks for one or more LBAs from the medium using the parameters specified in the Read-Write Error Recovery mode page (see 6.5.8)

3.1.68

read operation

process by which a device server performs operations as described in this standard

Note 1 to entry: Examples of read operations are read cache operations and read medium operations.

Note 2 to entry: See 4.2.3.

3.1.69

reassign

perform a reassign operation

3.1.70

reassign operation

operation during which the device server changes the assignment of an LBA from a specified physical block to another physical block and adds the specified physical block to the GLIST

3.1.71

recovered error

error for which a device server is able to read or write a logical block within the recovery limits specified in the Read-Write Error Recovery mode page (see 6.5.8) or the Verify Error Recovery mode page (see 6.5.9)

3.1.72

recovered read error

recovered error that occurs during a read medium operation

3.1.73

redundancy group

grouping of XOR-protected data (see 3.1.105) and associated check data (see 3.1.13) into a single type of data redundancy (see SCC-2)

3.1.74

resource provisioned logical unit

logical unit that may or may not store logical block data for every LBA and that provides enough LBA mapping resources (see 3.1.37) to map every LBA

Note 1 to entry: See 4.7.3.2.

3.1.75

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.

3.1.76**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-5.

3.1.77**sense data**

data describing command completed information that a device server delivers to an application client in the same I_T_L_Q nexus transaction as the status or as parameter data in response to a REQUEST SENSE command

Note 1 to entry: See SPC-4.

3.1.78**sense key**

contents of the SENSE KEY field in the sense data

Note 1 to entry: See SAM-5.

3.1.79**status**

one byte of response information that contains a coded value defined in SAM-5, transferred from a device server to an application client upon completion of each command

Note 1 to entry: See SAM-5.

3.1.80**stopped power condition**

power condition in which a device server terminates TEST UNIT READY commands and logical block access commands

Note 1 to entry: See 4.21.4.

3.1.81**synchronize cache operation**

process by which a device server synchronizes logical blocks within the volatile cache with the non-volatile cache or the medium

3.1.82**thin provisioned logical unit**

logical unit that may or may not store logical block data for every LBA and that may or may not provide enough LBA mapping resources (see 3.1.37) to map every LBA

Note 1 to entry: See 4.7.3.3.

3.1.83**threshold set**

set of two or more logical blocks used for tracking logical block provisioning thresholds

Note 1 to entry: See 4.7.3.7.

3.1.84**threshold set size**

number of LBAs in a threshold set

Note 1 to entry: See 4.7.3.7.

3.1.85

token

representation of a collection of data

Note 1 to entry: See SPC-4.

3.1.86

unit attention condition

state that a logical unit (see 3.1.43) 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 3.1.35)

Note 1 to entry: See SAM-5.

3.1.87

unmap command

command that requests an unmap operation

Note 1 to entry: See 4.2.2.

3.1.88

unmap operation

process by which a device server either deallocates or anchors a single LBA

Note 1 to entry: See 4.2.3 and 4.7.3.4.

3.1.89

unmapped

logical block provisioning state of an LBA (see 4.7.1) in which the LBA is either anchored or deallocated

3.1.90

unrecovered error

error for which a device server is unable to read a logical block or write a logical block within the recovery limits specified in the Read-Write Error Recovery mode page (see 6.5.8) and/or the Verify Error Recovery mode page (see 6.5.9)

3.1.91

unrecovered read error

unrecovered error that occurs during a read medium operation

3.1.92

unrecovered write error

unrecovered error that occurs during a write medium operation

3.1.93

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

3.1.94

user data segment

contiguous sequence of logical blocks

Note 1 to entry: See 4.28.

3.1.95

verify command

command that requests verify operations

Note 1 to entry: See 4.2.2.

3.1.96**verify medium operation**

process by which a device server reads logical blocks for one or more LBAs from the medium using the parameters specified in the Verify Error Recovery mode page (see 6.5.9)

3.1.97**verify operation**

process by which the device server performs operations as described in this standard

Note 1 to entry: An example of a verify operation is a verify medium operation.

Note 2 to entry: See 4.2.3.

3.1.98**volatile cache**

cache that does not retain logical block data between power cycles

3.1.99**volatile medium**

medium that does not retain logical block data between power cycles

Note 1 to entry: An example of volatile medium is a silicon memory device that loses data written to it if device power is lost.

3.1.100**write cache operation**

process by which a device server writes logical blocks for one or more LBAs to the cache (see 4.7.1)

3.1.101**write command**

command that requests write operations

Note 1 to entry: See 4.2.2.

3.1.102**write medium operation**

process by which a device server writes logical blocks for one or more LBAs to the medium using the parameters specified in the Read-Write Error Recovery mode page (see 6.5.8)

3.1.103**write operation**

process by which a device server performs operations as described in this standard

Note 1 to entry: Examples of write operations are write cache operations and write medium operations.

Note 2 to entry: See 4.2.3.

3.1.104**XOR operation**

processing of an XOR bitwise on two identical-sized multiple-bit input values

Note 1 to entry: An example of multiple-bit inputs is the current content of a logical block and the content contained in the Data-Out Buffer having the same number of bytes.

3.1.105**XOR-protected data**

logical blocks (i.e., including logical block data) that are part of a redundancy group

3.2 Symbols

Symbols used in this standard include:

Symbol	Meaning
+	plus
–	minus
×	multiplied by
÷	divided by
=	equals
≠	not equal to
<	less than
>	greater than

3.3 Abbreviations

Abbreviations used in this standard include:

Abbreviation	Meaning
CbCS	Capability based Command Security
CDB	command descriptor block (see 3.1.15)
CRC	cyclic redundancy check (see 3.1.17)
ECC	error correcting code (see 3.1.26)
FCP-4	Fibre Channel Protocol for SCSI, fourth version (see Bibliography)
GLIST	grown defect list (see 3.1.33)
LBA	logical block address (see 3.1.40)
LBP	logical block provisioning (see 4.7.4)
LSB	least significant bit
LUN	logical unit number
M	implementation is mandatory
MMC-6	SCSI Multimedia Commands - 6 (see Bibliography)
MSB	most significant bit
O	implementation is optional
PLIST	primary defect list (see 3.1.58)
n/a	not applicable
ROD	representation of data (see SPC-4)
SAM-5	SCSI Architecture Model - 5 (see clause 2)
SAS-3	Serial Attached SCSI-3 (see Bibliography)
SCSI	Small Computer System Interface family of standards
SCC-2	SCSI-3 Controller Commands - 2 (see clause 2)
SES-2	SCSI Enclosure Services - 2 (see clause 2)
SPC-4	SCSI Primary Commands - 4 (see clause 2)
SPL-2	SAS Protocol Layer-2 (see clause 2)
VPD	Vital product data (see 6.6)
XOR	exclusive-or (see 3.1.27)

3.4 Keywords

3.4.1

ignored

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

Note 1 to entry: The contents or value of an ignored bit, byte, word, field or code value shall not be examined by the receiving SCSI device and may be set to any value by the transmitting SCSI device.

3.4.2

invalid

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

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

3.4.3

mandatory

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

3.4.4

may

keyword that indicates flexibility of choice with no implied preference

Note 1 to entry: "May" is equivalent to "may or may not".

3.4.5

may not

keywords that indicate flexibility of choice with no implied preference

Note 1 to entry: "May not" is equivalent to "may or may not".

3.4.6

obsolete

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

3.4.7

optional

keyword that describes features that are not required to be implemented by this standard; however, if any optional feature defined in this standard is implemented, then it shall be implemented as defined in this standard

3.4.8

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

3.4.9

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. Receipt of reserved code values in defined fields shall be reported as an error.

3.4.10

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.

3.4.11

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.

3.4.12

should

keyword indicating flexibility of choice with a strongly preferred alternative

Note 1 to entry: "Should" is equivalent to the phrase "it is strongly recommended".

3.4.13

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.

3.5 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 this clause or in the text where they first appear.

Normal case is used for words having the normal English meaning.

Upper case is used when referring to names of commands, status codes, sense keys, and additional sense codes (e.g., REQUEST SENSE command).

If there is more than one CDB length for a particular command (e.g., ORWRITE (16) and ORWRITE (32)), and the name of the command is used in a sentence without any CDB length descriptor (e.g., ORWRITE), then the condition described in the sentence applies to all CDB lengths for that command.

Names of fields and state variables are in small uppercase (e.g. NAME). When a field or state variable name contains acronyms, uppercase letters may also be used for readability (e.g., the LOGERR bit). Normal case is used when the contents of a field or state variable are being discussed. Fields or state variables containing only one bit are usually referred to as the NAME bit instead of the NAME field.

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 listed 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).

In the event of conflicting information the precedence for requirements defined in this standard is:

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

Tables show data format and values. Not all tables or figures are fully described in the text.

Notes and examples do not constitute any requirements for implementers, and notes are numbered consecutively throughout this standard.

3.6 Numeric and character conventions

3.6.1 Numeric conventions

A binary number is represented in this standard by any sequence of digits comprised of only the Western-Arabic numerals 0 and 1 immediately followed by a lower-case b (e.g., 0101b). Underscores are included between characters in binary number representations to increase readability or delineate field boundaries (e.g., 0_0101_1010b).

A hexadecimal number is represented in this standard by any sequence of digits comprised of only the Western-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 are included in hexadecimal number representations to increase readability or delineate field boundaries (e.g., B_FD8C_FA23h).

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

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).

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 represented using various conventions.

Table 1 – Numbering convention examples

ISO/IEC	United States	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 $666\ \frac{2}{3}$, and $12.\overline{142\ 857}$ means $12.142\ 857\ 142\ 857\dots$ or $12\ \frac{1}{7}$).

3.6.2 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;
- 3) prefix symbol and unit:
 - 1) decimal prefix symbol (e.g., M) (see table 2); and
 - 2) unit abbreviation;

and

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

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

STANDARDSISO.COM · Click to view the full PDF of ISO/IEC 14776-323:2017

3.7 State machine conventions

Figure 1 shows how state machines are described in this standard.

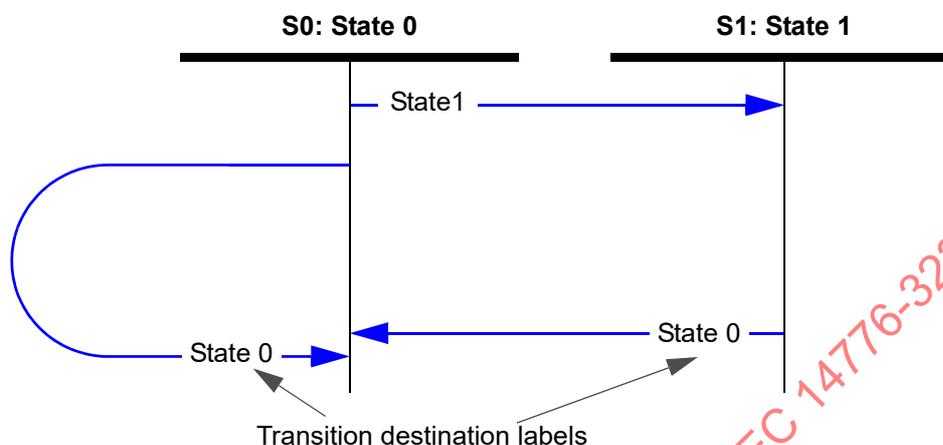


Figure 1 – Example state machine figure

The state machine figure is followed by subclauses describing the states and state transitions.

Each state and state transition is described in the list with particular attention to the conditions that cause the transition to occur and special conditions related to the transition.

A system specified in this manner has the following properties:

- a) time elapses only within discrete states; and
- b) state transitions are logically instantaneous.

4 Direct access block device type model

4.1 Direct access block device type model introduction

Table 3 shows the topics in clause 4 and a reference to the subclause where each topic is described.

Table 3 – Direct access block device type model topics

Topic	Reference
Direct access block device type model overview	4.2
Media examples	4.3
Removable media	4.4
Logical blocks	4.5
Physical blocks	4.6
Logical block provisioning	4.7
Data de-duplication	4.8
Ready state	4.9
Initialization	
Sanitize operations	4.11
Write protection	4.12
Medium defects	4.13
Write and unmap failures	4.14
Caches	4.15
Implicit HEAD OF QUEUE command processing	4.16
Reservations	4.17
Error reporting	4.18
Model for XOR commands	4.19
Rebuild assist mode	4.20
START STOP UNIT and power conditions	4.21
Protection information model	4.22
Grouping function	4.23
Background scan operations	4.24
Association between commands and CbCS permission bits	4.25
Deferred microcode activation	4.26
Model for uninterrupted sequences on LBA ranges	4.27
Referrals	4.28
ORWRITE commands	4.29
Block device ROD token operations	4.30

4.2 Direct access block device type model

4.2.1 Direct access block device type model overview

SCSI devices that conform to this standard are referred to as direct access block devices (e.g., hard disk drives, removable rigid disks, and solid state drives).

This standard is to be used in conjunction with SAM-5, SPC-4, SCC-2, and SES-2.

Direct access block devices store data in logical blocks for later retrieval.

Logical blocks are stored by a process that causes localized changes or transitions within a medium. The changes made to the medium to store the logical blocks may be volatile (i.e., not retained through power cycles) or non-volatile (i.e., retained through power cycles).

4.2.2 Logical block access command types

The following are logical block access command types:

- a) read commands;
- b) unmap commands;
- c) verify commands;
- d) write commands; and
- e) other commands (e.g., a FORMAT UNIT command or a SANITIZE command).

See table 31 for a list of commands for direct access block devices, including the logical block access command type. Some commands may be more than one type of logical block access command (e.g., a COMPARE AND WRITE command is both a read command and a write command).

4.2.3 Logical block access operation types

Each named command type (see 4.2.2) is processed by performing one or more of the following:

- a) read operations;
- b) unmap operations;
- c) verify operations; and
- d) write operations.

A device server that supports optional features (e.g., caches (see 4.15)) may be required to support additional requirements (e.g., cache coherency, LBA mapping resource allocations) that are related to those features for specific operations (e.g., read operations, write operations).

In a device server that does not support any optional features:

- a) any read operation causes only read medium operations to be performed;
- b) any verify operation causes only verify medium operations to be performed; and
- c) any write operation causes only write medium operations to be performed.

The requirements for any optional feature (e.g., caches) may include additional requirements for the operations described in this subclause.

If an optional feature (e.g., caches) defines requirements for read operations, then the device server shall support those requirements for both verify operations and read operations.

4.3 Media examples

4.3.1 Media examples overview

Examples of types of media used by the direct access block device are:

- a) a rotating medium (see 4.3.2); and
- b) a memory medium (see 4.3.3).

Other types of media are possible.

4.3.2 Rotating media

A rotating medium is one or more spinning disks, each coated with a magnetic material that allows flux changes to be induced and recorded. An actuator positions a read-write head radially across the spinning disk, allowing the device to randomly read or write the information at any radial position. Data is stored by using the write portion of the head to record flux changes and the recorded data is read by using the read portion of the head.

The circular path followed by the read-write head at a particular radius is called a track. A track is divided into sectors each containing blocks of stored data. If there is more than one disk spinning on a single axis and the actuator has a read-write head to access each of the disk surfaces, then the collection of tracks at a particular radius is called a cylinder.

A logical block is stored in one or more sectors, or a sector may store more than one logical block. Sectors may also contain information for accessing, synchronizing, and protecting the integrity of the logical blocks.

A rotating medium direct access block device is ready if:

- a) the disks are rotating at the correct speed; and
- b) the read-write circuitry is powered and ready to access the data.

A START STOP UNIT command (see 5.25) may be required to bring the logical unit to the ready state.

The rotating medium in a direct access block device is non-volatile.

4.3.3 Memory media

A memory medium is solid state, random access memory (RAM) (e.g., static RAM (SRAM), dynamic RAM (DRAM), magnetoresistive RAM (MRAM), ferroelectric RAM (FeRAM), or flash memory).

A memory medium direct access block device may be ready after power on and may not require a START STOP UNIT command (see 5.25) to bring the logical unit to a ready state.

These logical units may be nonmechanical, and logical blocks may be accessed with similar access times regardless of their location on the medium. Memory medium direct access block devices may store less data than disks or tapes and may be volatile.

A memory medium may be volatile (e.g., SRAM or DRAM) or non-volatile (e.g., SRAM or DRAM with battery backup, MRAM, FeRAM, or flash memory).

4.4 Removable media

The medium may be removable or non-removable. A removable medium may be contained within a cartridge or jacket to prevent damage to the recording surfaces.

A removable medium has an attribute of being mounted or demounted on a suitable transport mechanism in a direct access block device. A removable medium is mounted when the direct access block device is capable of accessing its medium (e.g., performing read medium operations and write medium operations). A removable medium is demounted at any other time (e.g., during loading, unloading, or storage).

An application client may check whether a removable medium is mounted by sending a TEST UNIT READY command (see SPC-4). A direct access block device containing a removable medium may not be accessible for read operations, unmap operations, and write operations until it receives a START STOP UNIT command with the START bit set to one (see 5.25).

If a direct access block device implements cache, either volatile or non-volatile, then the device server ensures that all logical blocks on the medium contain the most recent logical block data prior to permitting demounting of the removable medium.

If the medium in a direct access block device is removable, and the medium is removed, then the device server shall establish a unit attention condition with the additional sense code set to the appropriate value (e.g., MEDIUM NOT PRESENT). The PREVENT ALLOW MEDIUM REMOVAL command (see 5.10) allows an application client to restrict the demounting of the removable medium.

If an application client sends a START STOP UNIT command to request that the removable medium to be ejected and the direct access block device is prevented from demounting the medium by a previous PREVENT ALLOW MEDIUM REMOVAL command, then the START STOP UNIT command is terminated by the device server.

4.5 Logical Blocks

Logical blocks are stored on the medium. Logical blocks:

- a) contain logical block data that contains:
 - A) user data; and
 - B) protection information, if any;
 and
- b) may contain additional information (e.g., an ECC which may be used for medium defect management (see 4.13)), which may not be accessible to the application client.

The number of bytes of user data contained in each logical block is the logical block length. The logical block length is greater than or equal to one byte and should be an even number of bytes (e.g., 512 bytes, 520 bytes, 4 096 bytes, or 4 104 bytes). The logical block length does not include the length of protection information, if any, and additional information, if any, that are contained in the logical block. The logical block length is the same for all logical blocks in the logical unit. The LOGICAL BLOCK LENGTH field in the READ CAPACITY (10) parameter data (see 5.15.2) and the READ CAPACITY (16) parameter data (see 5.16.2) indicates the logical block length. The FORMAT UNIT command (see 5.3) and the mode parameter block descriptor (see 6.5.2) are used together by an application client to change the logical block length in direct access block devices that support changeable logical block lengths.

Each logical block is referenced by a unique LBA, which is represented as either four bytes in length or eight bytes in length. The LBAs on a logical unit shall begin with zero and shall be contiguous up to the last LBA on the logical unit. The last LBA is $[n-1]$, where $[n]$ is the number of logical blocks accessible by the application client. The RETURNED LOGICAL BLOCK ADDRESS field in the READ CAPACITY (10) parameter data (see 5.15.2) and the READ CAPACITY (16) parameter data (see 5.16.2) indicates the value of $[n-1]$. Each LBA has a logical block provisioning state (see 4.7) of mapped, deallocated, or anchored.

Some commands support only four-byte LOGICAL BLOCK ADDRESS fields (e.g., READ (10), and WRITE (10)).

If the capacity of the logical unit exceeds that accessible with four-byte LBAs, then the device server returns the RETURNED LOGICAL BLOCK ADDRESS field set to FFFF_FFFFh in the READ CAPACITY (10) parameter data, indicating that an application client should:

- a) enable descriptor format sense data (see SPC-4) in the Control mode page (see SPC-4) and in any REQUEST SENSE commands (see SPC-4) it sends; and
- b) use commands with eight-byte LOGICAL BLOCK ADDRESS fields (e.g., READ (16), and WRITE (16)).

NOTE 1 - If a command requests access to an LBA greater than FFFF_FFFFh, fixed format sense data is used, and an error occurs for that LBA, then there is no field in the sense data large enough to report that LBA as having an error (see 4.18).

If a command is received that references or attempts to access a logical block that exceeds the capacity of the medium, then the device server shall terminate the command (e.g., with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to LOGICAL BLOCK ADDRESS OUT OF RANGE). The device server:

- a) should terminate the command before processing; and
- b) may terminate the command after the device server has transferred some, all, or none of the data.

The location of a logical block on the medium is not required to have a relationship to the location of any other logical block. However, in a direct access block device with rotating media (see 4.3.2), the time to access a logical block at LBA [x+1] after accessing LBA [x] is often less than the time to access some other logical block.

4.6 Physical blocks

A physical block is a set of data bytes on the medium accessed by the device server as a unit. A physical block may contain:

- a) a portion of a logical block (i.e., there are multiple physical blocks in the logical block)(e.g., a physical block length of 512 bytes with a logical block length of 2 048 bytes);
- b) a single complete logical block; or
- c) more than one logical block (i.e., there are multiple logical blocks in the physical block)(e.g., a physical block length of 4 096 bytes with a logical block length of 512 bytes).

Each physical block may include additional information (e.g., an ECC which may be used for medium defect management (see 4.13)), which may not be accessible to the application client.

If the device server supports the COR_DIS bit and/or the WR_UNCOR bit in a WRITE LONG command (see 5.41 and 5.42), then the device server shall have the capability of marking individual logical blocks as containing pseudo unrecovered errors with correction enabled or with correction disabled.

Logical blocks may or may not be aligned to physical block boundaries. A mechanism for establishing the alignment is not defined by this standard.

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC 14776-323:2017

Figure 2 shows examples of where there are one or more physical blocks per logical block, and LBA 0 is aligned to a physical block boundary. The LOGICAL BLOCKS PER PHYSICAL BLOCK EXPONENT field and the LOWEST ALIGNED LOGICAL BLOCK ADDRESS field in the READ CAPACITY (16) parameter data (see 5.16.2) indicate the alignment.

The LOGICAL BLOCKS PER PHYSICAL BLOCK EXPONENT field is set to 0h (i.e., indicating one or more physical blocks per logical block).
 The LOWEST ALIGNED LOGICAL BLOCK ADDRESS field is set to 0000h (i.e., indicating that LBA 0 is located at the beginning of a physical block).

4 physical blocks per logical block:

LBA 0				LBA 1				...
PB	PB	PB	PB	PB	PB	PB	PB	...

3 physical blocks per logical block:

LBA 0			LBA 1			LBA 2			...
PB	PB	PB	PB	PB	PB	PB	PB	PB	...

2 physical blocks per logical block:

LBA 0		LBA 1		LBA 2		LBA 3		LBA 4		...
PB	PB	...								

1 physical block per logical block:

LBA 0	LBA 1	LBA 2	LBA 3	LBA 4	LBA 5	LBA 6	LBA 7	LBA 8	LBA 9	LBA 10	...
PB	...										

Key:

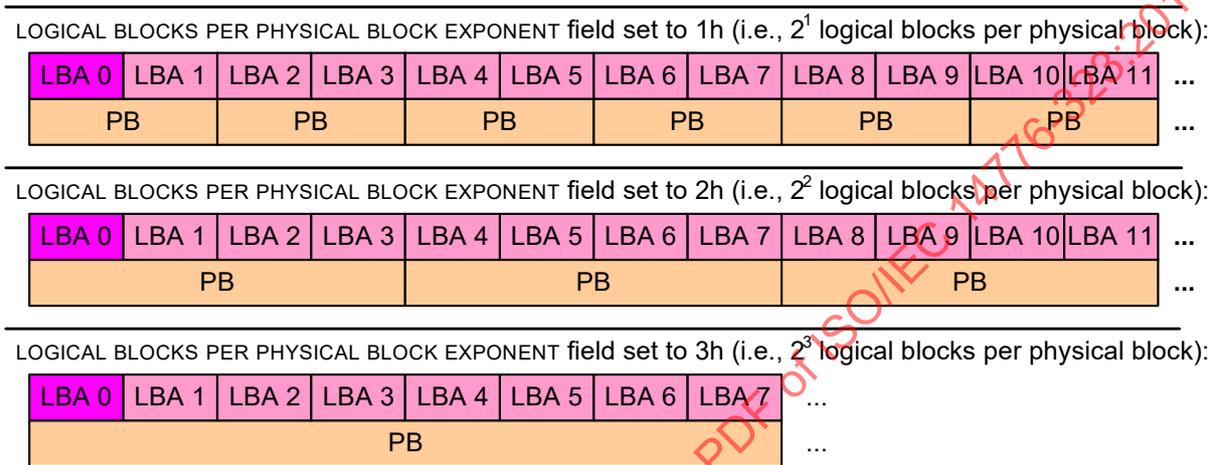
LBA n = logical block with LBA n
 PB = physical block

Figure 2 – One or more physical blocks per logical block examples

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC 14776-323:2017

Figure 3 shows examples of where there are one or more logical blocks per physical block, and LBA 0 is aligned to a physical block boundary. The LOGICAL BLOCKS PER PHYSICAL BLOCK EXPONENT field and the LOWEST ALIGNED LOGICAL BLOCK ADDRESS field in the READ CAPACITY (16) parameter data (see 5.16.2) indicate the alignment.

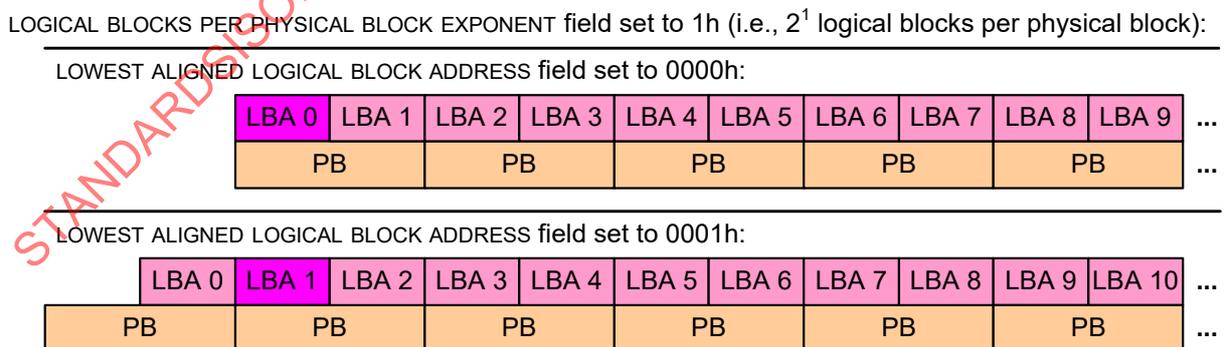
The LOGICAL BLOCKS PER PHYSICAL BLOCK EXPONENT field is set to a non-zero value (i.e., indicating more than one logical block per physical block).
 The LOWEST ALIGNED LOGICAL BLOCK ADDRESS field is set to 0000h (i.e., indicating that LBA 0 is located at the beginning of a physical block).



Key:
 LBA n = logical block with LBA n
 PB = physical block

Figure 3 – One or more logical blocks per physical block examples

Figure 4 shows examples of where there are two logical blocks per physical block, and different LBAs are aligned to physical block boundaries. The LOGICAL BLOCKS PER PHYSICAL BLOCK EXPONENT field and the LOWEST ALIGNED LOGICAL BLOCK ADDRESS field in the READ CAPACITY (16) parameter data (see 5.16.2) indicate the alignment.



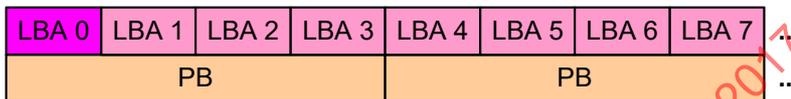
Key:
 LBA n = logical block with LBA n
 PB = physical block

Figure 4 – Two logical blocks per physical block alignment examples

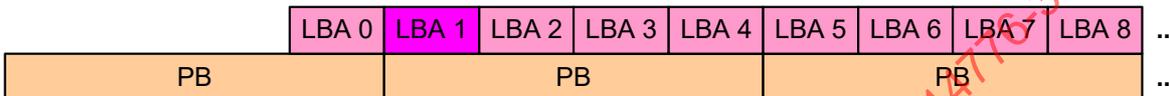
Figure 5 shows examples of where there are four logical blocks per physical block, and different LBAs are aligned to physical block boundaries. The LOGICAL BLOCKS PER PHYSICAL BLOCK EXPONENT field and the LOWEST ALIGNED LOGICAL BLOCK ADDRESS field in the READ CAPACITY (16) parameter data (see 5.16.2) indicate the alignment.

LOGICAL BLOCKS PER PHYSICAL BLOCK EXPONENT field set to 2h (i.e., 2² logical blocks per physical block):

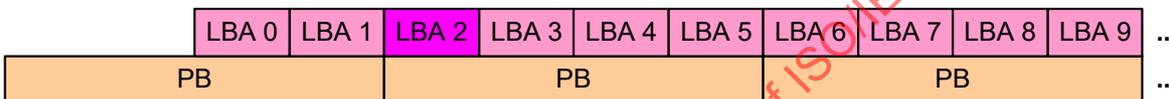
LOWEST ALIGNED LOGICAL BLOCK ADDRESS field set to 0000h:



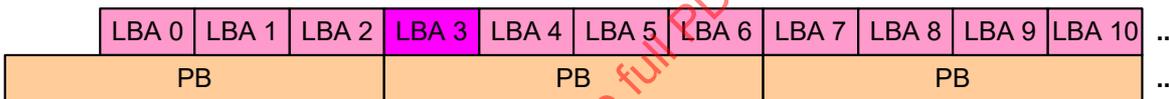
LOWEST ALIGNED LOGICAL BLOCK ADDRESS field set to 0001h:



LOWEST ALIGNED LOGICAL BLOCK ADDRESS field set to 0002h:



LOWEST ALIGNED LOGICAL BLOCK ADDRESS field set to 0003h:



Key:

LBA n = logical block with LBA n
 PB = physical block

Figure 5 – Four logical blocks per physical block alignment examples

If there is more than one logical block per physical block, then not all of the logical blocks are aligned to the physical block boundaries. When using logical block access commands (see 4.2.2), application clients should:

- a) specify an LBA that is aligned to a physical block boundary; and
- b) access an integral number of physical blocks, provided that the access does not go beyond the last LBA on the medium.

See annex E for an example method in which application clients may use alignment information to determine optimal performance for logical block access.

4.7 Logical block provisioning

4.7.1 Logical block provisioning overview

Each LBA in a logical unit is either mapped or unmapped. For LBAs that are mapped, there is a known relationship between the LBA and one or more physical blocks that contain logical block data. For LBAs that are unmapped, the relationship between the LBA and a physical block is not defined. Figure 6 shows two examples of the relationship between mapped and unmapped LBAs and physical blocks in a logical unit. One example shows one LBA per physical block and one example shows two LBAs per physical block. The LOGICAL BLOCKS PER PHYSICAL BLOCK EXPONENT field is defined in the READ CAPACITY (16) parameter data (see 5.16.2).

LOGICAL BLOCKS PER PHYSICAL BLOCK EXPONENT field set to 0h (i.e., 2^0 logical blocks per physical block):

LBA 0	LBA 1	LBA 2	LBA 3	LBA 4	LBA 5	LBA 6	LBA 7
PB	PB	Unmapped	PB	Unmapped	Unmapped	PB	PB

LOGICAL BLOCKS PER PHYSICAL BLOCK EXPONENT field set to 1h (i.e., 2^1 logical blocks per physical block):

LBA 0	LBA 1	LBA 2	LBA 3	LBA 4	LBA 5	LBA 6	LBA 7
PB		Unmapped		PB		Unmapped	

Key:

LBA n = logical block with LBA n

PB = physical block

Unmapped = the relationship between the LBA(s) and a physical block is not defined

Figure 6 – Examples of the relationship between mapped and unmapped LBAs and physical blocks

Each unmapped LBA is either anchored or deallocated. Anchored and deallocated are states in the LBP state machine (see 4.7.4) that have the following properties:

- a) a write command that specifies an anchored LBA does not require allocation of additional LBA mapping resources for that LBA; and
- b) a write command that specifies a deallocated LBA may require allocation of LBA mapping resources.

Depending on the logical block provisioning types (see table 4), the quantity of LBA mapping resources available to a logical unit may be greater than, equal to, or less than the quantity required to store logical block data for every LBA.

Table 4 list the logical block provisioning states supported by each type of logical block provisioning.

Table 4 – Logical block provisioning states supported by logical block provisioning type

Type	Logical block provisioning states			Reference
	Mapped	Unmapped		
		Anchored	Deallocated	
Fully provisioned	Mandatory	Prohibited	Prohibited	4.7.2
Resource provisioned	Mandatory	Mandatory	Prohibited	4.7.3.2
Thin provisioned	Mandatory	Optional	Mandatory	4.7.3.3

4.7.2 Fully provisioned logical unit

The device server shall map every LBA in a fully provisioned logical unit. A fully provisioned logical unit shall provide enough LBA mapping resources to contain all logical blocks for the logical unit's capacity as reported in the READ CAPACITY (10) parameter data (see 5.15.2) and the READ CAPACITY (16) parameter data (see 5.16.2). The device server shall not cause any LBA on a fully provisioned logical unit to become unmapped (i.e., anchored or deallocated).

A fully provisioned logical unit does not support logical block provisioning management (see 4.7.3). A fully provisioned logical unit may support the GET LBA STATUS command (see 5.4).

The device server in a fully provisioned logical unit shall set the LBPME bit to zero in the READ CAPACITY (16) parameter data (see 5.16.2).

4.7.3 Logical block provisioning management

4.7.3.1 Logical block provisioning management overview

A logical unit that supports logical block provisioning management (i.e., implements unmapped LBAs, unmap operations, and related actions) shall be either:

- a) resource provisioned (see 4.7.3.2); or
- b) thin provisioned (see 4.7.3.3).

A logical unit that supports logical block provisioning management shall implement the LBP state machine (see 4.7.4) for each LBA.

The device server in a logical unit that supports logical block provisioning management:

- a) shall support the Logical Block Provisioning VPD page (see 6.6.4);
- b) may supply a provisioning group designation descriptor as defined in the Logical Block Provisioning VPD page;
- c) may support logical block provisioning thresholds (see 4.7.3.7.1);
- d) may support the GET LBA STATUS command (see 5.4);
- e) should support the Block Limits VPD page (see 6.6.3); and
- f) shall support at least one of the following unmap mechanisms:
 - A) the UNMAP command (see 5.28);
 - B) the UNMAP bit in the WRITE SAME (10) command (see 5.43);
 - C) the UNMAP bit in the WRITE SAME (16) command (see 5.44); or
 - D) the UNMAP bit in the WRITE SAME (32) command (see 5.45).

If the device server supports:

- a) the UNMAP bit in the WRITE SAME (10) command or in the WRITE SAME (16) command; and
- b) the WRITE SAME (32) command (see 5.45),

then the device server shall support the UNMAP bit in the WRITE SAME (32) command.

If a device server supports the UNMAP command and the Block Limits VPD page, then the device server shall:

- a) set the MAXIMUM UNMAP LBA COUNT field in the Block Limits VPD page to a value greater than or equal to one; and
- b) set the MAXIMUM UNMAP DESCRIPTOR COUNT field in the Block Limits VPD page to a value greater than or equal to one.

4.7.3.2 Resource provisioned logical unit

A resource provisioned logical unit shall support logical block provisioning management (see 4.7.3.1).

The device server in a resource provisioned logical unit assigns mapping resources to every LBA for the capacity indicated by the RETURNED LOGICAL BLOCK ADDRESS field in the READ CAPACITY (10) parameter data (see 5.15.2) and the READ CAPACITY (16) parameter data (see 5.16.2).

The device server shall map or anchor (i.e., not deallocate) each LBA in a resource provisioned logical unit. A resource provisioned logical unit shall provide LBA mapping resources sufficient to map all LBAs for the logical unit's capacity as indicated in the RETURNED LOGICAL BLOCK ADDRESS field of the READ CAPACITY (10) parameter data (see 5.15.2) and the READ CAPACITY (16) parameter data (see 5.16.2). A resource provisioned logical unit may provide resources in excess of this requirement. The device server shall not cause any LBA on a resource provisioned logical unit to become deallocated.

The device server in a resource provisioned logical unit shall set:

- a) the LBPME bit to one in the READ CAPACITY (16) parameter data (see 5.16.2);
- b) the PROVISIONING TYPE field to 001b (i.e., resource provisioned) in the Logical Block Provisioning VPD page (see 6.6.4); and
- c) the ANC_SUP bit to one in the Logical Block Provisioning VPD page.

The initial condition of every LBA in a resource provisioned logical unit is anchored (see 4.7.4.1).

4.7.3.3 Thin provisioned logical unit

A thin provisioned logical unit shall support logical block provisioning management (see 4.7.3.1).

The device server in a thin provisioned logical unit may indicate a larger capacity in the RETURNED LOGICAL BLOCK ADDRESS field in the READ CAPACITY (10) parameter data (see 5.15.2) and the READ CAPACITY (16) parameter data (see 5.16.2) than the number of LBA mapping resources available for mapping LBAs in the logical unit.

The device server shall map, anchor, or deallocate each LBA in a thin provisioned logical unit (see table 4). A thin provisioned logical unit is not required to provide LBA mapping resources sufficient to map all LBAs for the logical unit's capacity as indicated in the RETURNED LOGICAL BLOCK ADDRESS field of the READ CAPACITY (10) parameter data (see 5.15.2) and the READ CAPACITY (16) parameter data (see 5.16.2).

If the logical unit does not support anchored LBAs (i.e., the ANC_SUP bit is set to zero in the Logical Block Provisioning VPD page (see 6.6.4)), then:

- a) every unmapped LBA in the logical unit shall be deallocated; and
- b) the device server shall terminate every command that specifies anchoring an LBA (e.g., a WRITE SAME command with the ANCHOR bit set to one (see 5.43)).

The device server in a thin provisioned logical unit shall set:

- a) the LBPME bit to one in the READ CAPACITY (16) parameter data (see 5.16.2); and
- b) the PROVISIONING TYPE field to 010b (i.e., thin provisioned) in the Logical Block Provisioning VPD page (see 6.6.4).

The initial condition of every LBA in a thin provisioned logical unit is deallocated (see 4.7.4.1).

4.7.3.4 Unmapping LBAs

4.7.3.4.1 Processing unmap requests

Application clients use unmap commands (see table 31) to request that LBAs be unmapped. For each LBA that is requested to be unmapped, the device server shall:

- a) perform an unmap operation (see 4.7.3.4.2) on the LBA; or
- b) make no change to the logical block provisioning state of the LBA.

The application client determines the logical block provisioning state of LBAs using the GET LBA STATUS command (see 5.4).

4.7.3.4.2 Unmap operations

An unmap operation:

- a) results in a single LBA becoming either deallocated or anchored;
- b) may change the relationship between one LBA and one or more physical blocks; and

- c) may change the logical block data that is returned in response to a subsequent read command specifying that LBA.

The data in all other mapped LBAs on the medium shall be preserved. Performing an unmap operation on an unmapped LBA shall not be considered an error.

An unmap operation may or may not release LBA mapping resources.

An application client may use an unmap command (see 4.2.2) to request that the device server perform one or more unmap operations. A single unmap command may result in zero or more unmap operations.

4.7.3.4.3 WRITE SAME command and unmap operations

A WRITE SAME command (see 5.43, 5.44, and 5.45) may be used to request unmap operations that deallocate or anchor the specified LBAs. If unmap operations are requested in a WRITE SAME command, then for each specified LBA:

- a) if the Data-Out Buffer of the WRITE SAME command is the same as the logical block data returned by a read operation from that LBA while in the unmapped state (see 4.7.4.5), then:
 - 1) the device server performs the actions described in table 5; and
 - 2) if an unmap operation is not performed in step 1), then the device server shall perform the specified write operation to that LBA;
- or
- b) if the Data-Out Buffer of the WRITE SAME command is not the same as the logical block data returned by a read operation from that LBA while in the unmapped state (see 4.7.4.5), then the device server shall perform the specified write operation to that LBA.

Table 5 – WRITE SAME command and unmap operations

Logical block provisioning management type	Unmap operations that request to deallocate the specified LBA	Unmap operations that request to anchor the specified LBA
Thin provisioned logical unit	a) should perform an unmap operation to deallocate the LBA (see 4.7.4.7.1); or b) may perform an unmap operation to anchor the LBA (see 4.7.4.8.1)	should perform an unmap operation to anchor the LBA
Resource provisioned logical unit	should perform an unmap operation to anchor the LBA	should perform an unmap operation to anchor the LBA

A WRITE SAME command shall not cause an LBA to become unmapped if unmapping that LBA creates a case in which a subsequent read of that unmapped LBA is able to return logical block data that differs from the Data-Out Buffer for that WRITE SAME command (see 4.7.4.5).

If the device server does not support allowing a WRITE SAME command to request unmap operations, then the device server shall:

- a) perform the write operations specified by the WRITE SAME command; and
- b) not perform any unmap operations.

The device server shall perform the write operations specified by a WRITE SAME command and shall not perform any unmap operations if the device server sets the LBPRZ bit to one in the READ CAPACITY (16) parameter data (see 5.16.2), and:

- a) any bit in the user data transferred from the Data-Out Buffer is not zero; or
- b) the protection information, if any, transferred from the Data-Out Buffer is not set to FFFF_FFFF_FFFF_FFFFh.

4.7.3.5 Autonomous LBA transitions

A device server may perform the following actions at any time:

- a) transition any deallocated LBA to mapped;
- b) transition any anchored LBA to mapped; or
- c) transition any deallocated LBA to anchored.

If the LBPRZ bit in the READ CAPACITY (16) parameter data (see 5.16.2) is set to one, and a mapped LBA references a logical block that contains:

- a) user data with all bits set to zero; and
- b) protection information, if any, set to FFFF_FFFF_FFFF_FFFFh,

then the device server may transition that mapped LBA to anchored or deallocated at any time.

The logical block provisioning state machine (see 4.7.4) specifies additional requirements for the transitions specified in this subclause.

4.7.3.6 Thin provisioned logical unit resource exhaustion considerations

If:

- a) a write operation is requested by an application client, and a temporary lack of LBA mapping resources prevents the logical unit from performing the write operation; or
- b) an unmap operation that transitions an LBA to the anchored state is requested by an application client and a temporary lack of LBA mapping resources prevents the logical unit from anchoring the LBA,

then the device server shall terminate the command requesting the operation with CHECK CONDITION status with the sense key set to NOT READY and the additional sense code set to LOGICAL UNIT NOT READY, SPACE ALLOCATION IN PROGRESS. In this case, the application client should resend the command.

If:

- a) a write operation is requested by an application client, and a persistent lack of LBA mapping resources prevents the logical unit from performing the write operation; or
- b) an unmap operation that transitions an LBA to the anchored state is requested by an application client and a persistent lack of LBA mapping resources prevents the logical unit from anchoring the LBA,

then the device server shall terminate the command requesting the unmap operation with CHECK CONDITION status with the sense key set to DATA PROTECT and the additional sense code set to SPACE ALLOCATION FAILED WRITE PROTECT. This condition shall not cause the device server to set the WP bit in the DEVICE-SPECIFIC PARAMETER field of the mode page header to one (see 6.5.1). In this case, recovery actions by the application client are outside the scope of this standard.

A logical block provisioning threshold may be available to monitor the availability of LBA mapping resources (see 4.7.3.7). A logical block provisioning log parameter that reports available LBA mapping resources may be available in the Logical Block Provisioning log page (see 6.4.4).

4.7.3.7 Logical block provisioning thresholds

4.7.3.7.1 Logical block provisioning thresholds overview

Logical block provisioning thresholds provide a mechanism for the device server to establish a unit attention condition to notify application clients when thresholds related to logical block provisioning are crossed. Logical block provisioning thresholds may operate on an armed increasing basis or an armed decreasing basis.

If a device server supports logical block provisioning thresholds, then the device server:

- a) shall support the Logical Block Provisioning mode page (see 6.5.7); and
- b) may support the Logical Block Provisioning log page (see 6.4.4).

The end points of the range over which a logical block provisioning threshold operates are defined as follows:

$$\text{threshold minimum} = ((\text{threshold count} \times \text{threshold set size}) - (\text{threshold set size} \times 0.5))$$

$$\text{threshold maximum} = ((\text{threshold count} \times \text{threshold set size}) + (\text{threshold set size} \times 0.5))$$

where:

- threshold minimum is the lowest number of LBAs in the range for this threshold;
- threshold maximum is the highest number of LBAs in the range for this threshold;
- threshold count is the center of the threshold range for this threshold (i.e., the threshold count value as specified in the threshold descriptor in the Logical Block Provisioning mode page); and
- threshold set size is the number of LBAs in each threshold set (i.e., $2^{(\text{threshold exponent})}$ LBAs where the threshold exponent is indicated in the Logical Block Provisioning VPD page (see 6.6.4)).

Table 6 defines the meaning of the combinations of values for the THRESHOLD RESOURCE field, the THRESHOLD TYPE field, and the THRESHOLD ARMING field that are used for logical block provisioning thresholds. See the Logical Block Provisioning mode page (see 6.5.7) for the definition of these fields.

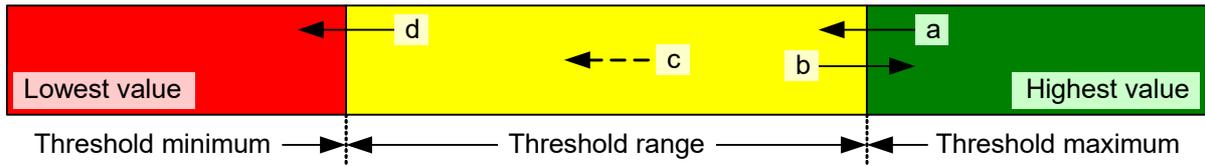
Table 6 – Threshold resource value, threshold type value, and threshold arming value for logical block provisioning thresholds

Threshold resource value	Threshold type value	Threshold arming value	Description
01h	000b	000b	The device server applies the threshold to the availability of LBA mapping resources and performs notifications as the availability of those resources decreases. ^a
02h	000b	001b	The device server applies the threshold to the usage of LBA mapping resources and performs notifications as the usage of those resources increases.
All other combinations			Reserved
^a The point when availability of LBA mapping resources reaches zero corresponds to the persistent lack of LBA mapping resources described in 4.7.3.6.			

4.7.3.7.2 Logical block provisioning armed decreasing thresholds

Figure 7 shows the operation of a logical block provisioning armed decreasing threshold. Figure 7 represents the entire range of possible values over which the threshold is being applied (e.g., for an available resource, the lowest value represents zero available resources and the highest value represents the maximum possible number of available resources).

If enabled, reporting of armed decreasing threshold events (i.e., the THRESHOLD ARMING field is set to 000b in the threshold descriptor in the Logical Block Provisioning mode page (see 6.5.7)) operates as shown in figure 7.



Notes:

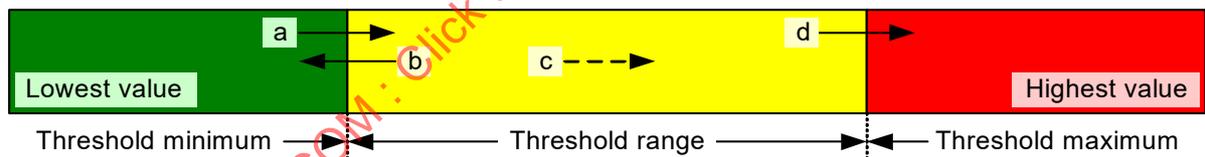
- a) if the value to which the threshold is being applied drops below the threshold maximum for the threshold range, then the notification trigger shall be enabled;
- b) if the value to which the threshold is being applied increases above the threshold maximum for the threshold range, then the notification trigger shall be disabled;
- c) if the notification trigger is enabled, then the device server may disable the notification trigger and perform logical block provisioning threshold notification (see 4.7.3.7.4); and
- d) if the notification trigger is enabled and the value to which the threshold is being applied drops below the threshold minimum for the threshold range, then the device server shall disable the notification trigger and perform logical block provisioning threshold notification as defined in 4.7.3.7.4.

Figure 7 – Armed decreasing threshold operation

4.7.3.7.3 Logical block provisioning armed increasing thresholds

Figure 8 shows the operation of a logical block provisioning armed increasing threshold. Figure 8 represents the entire range of possible values over which the threshold is being applied (e.g., for tracking usage of a resource, the lowest value represents zero resources being used and the highest value represents the maximum possible number of resources being used).

If enabled, reporting of armed increasing threshold events (i.e., the THRESHOLD ARMING field is set to 001b in the threshold descriptor in the Logical Block Provisioning mode page (see 6.5.7)) operates as shown in figure 8.



Notes:

- a) if the value to which the threshold is being applied increases above the threshold minimum for the threshold range, then the notification trigger shall be enabled;
- b) if the value to which the threshold is being applied decreases below the threshold minimum for the threshold range, then the notification trigger shall be disabled;
- c) if the notification trigger is enabled, then the device server may disable the notification trigger and perform logical block provisioning threshold notification (see 4.7.3.7.4); and
- d) if the notification trigger is enabled and the value to which the threshold is being applied increases above the threshold maximum for the threshold range, then the device server shall disable the notification trigger and perform logical block provisioning threshold notification as defined in 4.7.3.7.4.

Figure 8 – Armed increasing threshold operation

4.7.3.7.4 Logical block provisioning threshold notification

If the LBPERE bit is set to one in the Read-Write Error Recovery mode page (see 6.5.8), then logical block provisioning threshold notification is enabled and the device server shall perform notification for thresholds with the THRESHOLD TYPE field set to 000b in the threshold descriptor in the Logical Block Provisioning mode page (see 6.5.7) as follows:

- a) if the SITUA bit is set to one in the Logical Block Provisioning mode page, then:
 - A) if the device server has not established a unit attention condition as a result of this threshold being crossed since the last logical unit reset (see SAM-5) and a command through which the device server is able to report a unit attention condition arrives on any I_T nexus, then the device server shall establish a unit attention condition with the additional sense code set to THIN PROVISIONING SOFT THRESHOLD REACHED for only the SCSI initiator port associated with the I_T nexus on which that command was received before processing that command; or
 - B) if the device server has established a unit attention condition as a result of this threshold being crossed since the last logical unit reset and a command through which the device server is able to report a unit attention condition arrives on any I_T nexus, then the device server should establish a unit attention condition with the additional sense code set to THIN PROVISIONING SOFT THRESHOLD REACHED for only the SCSI initiator port associated with the I_T nexus on which that command was received before processing that command unless establishment of the unit attention condition causes a vendor specific frequency of unit attention conditions for this threshold to be exceeded;
- or
- b) if the SITUA bit is set to zero, then:
 - A) if the device server has not established a unit attention condition for the SCSI initiator port associated with all I_T nexuses as a result of this threshold being crossed since the last logical unit reset (see SAM-5), then the device server shall establish a unit attention condition with the additional sense code set to THIN PROVISIONING SOFT THRESHOLD REACHED for the SCSI initiator port associated with every I_T nexus; or
 - B) if the device server has established a unit attention condition for the SCSI initiator ports associated with all I_T nexuses as a result of this threshold being crossed since the last logical unit reset, then the device server should establish a unit attention condition with the additional sense code set to THIN PROVISIONING SOFT THRESHOLD REACHED for the SCSI initiator port associated with every I_T nexus, unless establishment of the unit attention condition causes a vendor specific frequency of unit attention conditions for this threshold to be exceeded.

If a unit attention condition is established as described in this subclause, then the device server shall report the following value in the INFORMATION field in the sense data (see SPC-4):

- a) the byte offset in the Logical Block Provisioning mode page of the first byte of the threshold descriptor to which this threshold notification applies.

If a unit attention condition with the additional sense code set to THIN PROVISIONING SOFT THRESHOLD REACHED is received by the application client, then the application client should reissue the command and take further recovery actions (e.g., administrator notification or other administrator actions). These recovery actions are outside the scope of this standard.

If the LBPERE bit is set to zero, then logical block provisioning threshold notification is disabled and the device server shall not establish any unit attention condition with the additional sense code set to THIN PROVISIONING SOFT THRESHOLD REACHED.

An additional sense code set to THIN PROVISIONING SOFT THRESHOLD REACHED is applicable to both thin provisioned logical units (see 4.7.3.3) and resource provisioned logical units (see 4.7.3.2).

4.7.4 LBP (logical block provisioning) state machine

4.7.4.1 LBP state machine overview

The LBP (logical block provisioning) state machine describes the mapping and unmapping of a single LBA by the device server for a thin provisioned logical unit (see 4.7.3.3) or a resource provisioned logical unit (see 4.7.3.2). This state machine does not apply to fully provisioned logical units (see 4.7.2).

There is one instance of this state machine for each LBA. If a command requests mapping or unmapping of more than one LBA, then there may be an independent transition in each instance of the state machine (e.g., each LBA may individually transition from mapped to deallocated or from anchored to deallocated).

4.7.4.2 LBP state machine for thin provisioned logical units supporting anchored LBAs

If the logical unit supports anchored LBAs (i.e., the ANC_SUP bit is set to one) and deallocated LBAs (i.e., is a thin provisioned logical unit), then this state machine consists of the following states:

- a) LBP1:Mapped state (see 4.7.4.6);
- b) LBP2:Deallocated state (see 4.7.4.7) (initial state); and
- c) LBP3:Anchored state (see 4.7.4.8).

The initial state of the state machine associated with each LBA is the LBP2:Deallocated state.

Figure 9 describes the LBP state machine for a logical unit that supports anchored LBAs and deallocated LBAs.

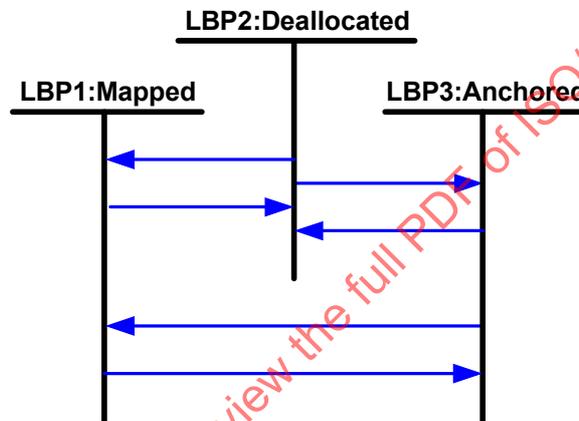


Figure 9 – LBP state machine (anchored LBAs supported and deallocated LBAs supported)

4.7.4.3 LBP state machine for thin provisioned logical units not supporting anchored LBAs

If the logical unit does not support anchored LBAs (i.e., is a thin provisioned logical unit and the ANC_SUP bit is set to zero), then this state machine consists of the following states:

- a) LBP1:Mapped state(see 4.7.4.6); and
- b) LBP2:Deallocated state (see 4.7.4.7) (initial state).

The initial state of the state machine associated with each LBA is the LBP2:Deallocated state.

Figure 10 describes the LBP state machine for a logical unit that does not support anchored LBAs.

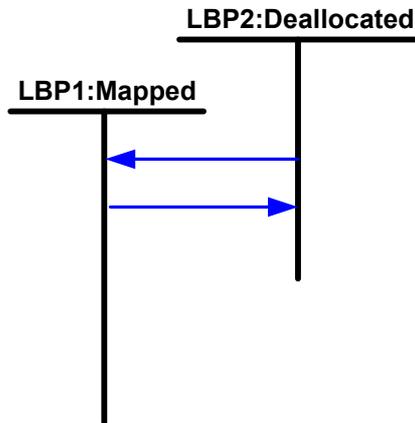


Figure 10 – LBP state machine (anchored LBAs not supported)

4.7.4.4 LBP state machine for resource provisioned logical units

If the logical unit does not support deallocated LBAs (i.e., is a resource provisioned logical unit and the ANC_SUP bit is set to one), then this state machine consists of the following states:

- a) LBP1:Mapped state (see 4.7.4.6)); and
- b) LBP3:Anchored state (see 4.7.4.8) (initial state).

The initial state of the state machine associated with each LBA is the LBP3:Anchored state for a resource provisioned logical unit (see 4.7.3.2).

Figure 11 describes the LBP state machine for a logical unit that does not support deallocated LBAs.

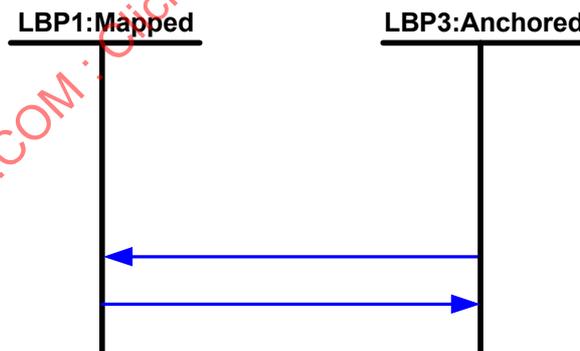


Figure 11 – LBP state machine (deallocated LBAs not supported)

4.7.4.5 Performing read operations with respect to logical block provisioning

Table 7 defines the logical block data that a read operation shall return for a mapped LBA.

Table 7 – Logical block data returned by a read operation from a mapped LBA

Condition	Logical block data returned
The LBA became mapped as the result of a format operation or sanitize operation and no write command has specified that LBA since the LBA became mapped.	The logical block data that was written to that LBA by the format operation or the sanitize operation.
The LBA became mapped as the result of a write command and no additional write command has specified that LBA since the LBA was mapped.	The logical block data that was written to that LBA by that write command.
The LBA became mapped as the result of an autonomous transition, and no write command has specified that LBA since the LBA was mapped.	The logical block data that would be returned if that autonomous transition had not occurred and the LBA had remained unmapped (see table 8).
A write command has specified that LBA since that LBA was mapped.	The logical block data that was most recently written to that LBA.

Table 8 defines the logical block data that a read operation shall return for an unmapped LBA.

Table 8 – Logical block data returned by a read operation from an unmapped LBA

LBPRZ bit ^a	Method used to unmap the LBA	Logical block data returned
0	see ^b	The value specified to be written if the write operation had been performed.
	see ^c	Logical block data containing: a) user data set to a vendor specific value that is not obtained from any other LBA, and b) protection information, if any, set to FFFF_FFFF_FFFF_FFFFh.
1	Any	Logical block data containing: a) user data set to zero; and b) protection information, if any, set to FFFF_FFFF_FFFF_FFFFh.

^a The LBPRZ bit is in the READ CAPACITY (16) parameter data (see 5.16.2).
^b A command for which the device server is allowed to perform either:
a) a write with specified logical block data; or
b) an unmap operation if the logical block data returned by read operations from unmapped LBAs matches the logical block data specified for the command that resulted in the unmap operation.
These commands are:
a) a FORMAT UNIT command specifying an initialization pattern;
b) a SANITIZE command specifying a sanitize overwrite operation; and
c) a WRITE SAME command with the UNMAP bit set to one.
^c These methods include but are not limited to:
a) a FORMAT UNIT command not specifying an initialization pattern;
b) a REASSIGN BLOCKS command;
c) a SANITIZE command specifying a sanitize block erase operation or a sanitize cryptographic erase operation; and
d) an UNMAP command.

After a read operation returns a value for an LBA, subsequent read operations from that LBA shall return the same value until a subsequent command alters the logical block data in that LBA (e.g., a write command or an unmap command (see table 31)).

4.7.4.6 LBP1:Mapped state**4.7.4.6.1 LBP1:Mapped state description**

Upon entry into this state, the relationship between the LBA and the physical block(s) that contains the logical block for that LBA shall be established.

If this state was entered from the LBP2:Deallocated state (see 4.7.4.7), then the device server shall allocate LBA mapping resources, if any, required to map the LBA.

If this state was entered from the LBP3:Anchored state (see 4.7.4.8), then:

- a) the device server shall not allocate LBA mapping resources; and
- b) the resource exhaustion conditions described in 4.7.3.6 shall not occur.

4.7.4.6.2 Transition LBP1:Mapped to LBP2:Deallocated

This transition shall occur after:

- a) an unmap operation that results in the deallocation of the LBA that was mapped.

This transition may occur at any time if the LBPRZ bit is set to one in the READ CAPACITY (16) parameter data (see 5.16.2), and the mapped LBA references a logical block that contains:

- a) user data with all bits set to zero; and
- b) protection information, if any, set to FFFF_FFFF_FFFF_FFFFh.

4.7.4.6.3 Transition LBP1:Mapped to LBP3:Anchored

This transition shall occur after:

- a) an unmap operation that results in the anchoring of the LBA that was mapped.

This transition may occur at any time if the LBPRZ bit is set to one in the READ CAPACITY (16) parameter data (see 5.16.2), and the mapped LBA references a logical block that contains:

- a) user data with all bits set to zero; and
- b) protection information, if any, set to FFFF_FFFF_FFFF_FFFFh.

4.7.4.7 LBP2:Deallocated state**4.7.4.7.1 LBP2:Deallocated state description**

While in this state:

- a) there shall be no relationship between the LBA and any physical block(s); and
- b) the device server should deallocate LBA mapping resources after they are no longer in use.

For an LBA in this state, an unmap operation that specifies deallocation of the LBA shall not cause a transition from this state.

The device server shall process a read command specifying an LBA in this state (i.e., a deallocated LBA) as described in table 8.

4.7.4.7.2 Transition LBP2:Deallocated to LBP1:Mapped

This transition:

- a) shall occur after a write operation to the LBA that was deallocated; or
- b) may occur at any time for reasons outside the scope of this standard.

4.7.4.7.3 Transition LBP2:Deallocated to LBP3:Anchored

This transition:

- a) shall occur after an unmap operation that results in anchoring of the LBA that was deallocated; or
- b) may occur at any time for reasons outside the scope of this standard.

4.7.4.8 LBP3:Anchored state

4.7.4.8.1 LBP3:Anchored state description

Upon entry into this state:

- a) LBA mapping resources shall be associated with the LBA; and
- b) there may or may not be a relationship between the LBA and physical block(s).

If this state was entered from the LBP2:Deallocated state, then the device server shall allocate LBA mapping resources, if any, required to anchor the LBA.

If this state was entered from the LBP1:Mapped state, then:

- a) the device server shall not allocate LBA mapping resources;
- b) the device server relies on LBA mapping resources already allocated to the LBA; and
- c) the resource exhaustion conditions described in 4.7.3.6 shall not occur.

For an LBA in this state, an unmap operation that specifies anchoring of the LBA shall not cause a transition from of this state.

The device server shall process a read command specifying an LBA in this state (i.e., an anchored LBA) as described in table 8.

4.7.4.8.2 Transition LBP3:Anchored to LBP1:Mapped

This transition:

- a) shall occur after a write operation to the LBA that was anchored; or
- b) may occur at any time for reasons outside the scope of this standard.

4.7.4.8.3 Transition LBP3:Anchored to LBP2:Deallocated

This transition shall occur after:

- a) an unmap operation that results in the deallocation of the LBA that was anchored.

4.8 Data de-duplication

Data de-duplication is the ability of a device server to recognize redundant or duplicate data and reduce the number of duplicate or redundant copies of the data while maintaining the application client supplied LBAs of the duplicate or redundant copies of the data. De-duplication shall not affect protection information, if any. Logical units that support data de-duplication may report a count of LBA resources that have been made available as a result of being de-duplicated (see 6.4.4.4).

4.9 Ready state

A direct access block device is ready when the device server is capable of processing logical block access commands that require access to the medium (see 4.2.2).

A direct access block device using a removable medium (see 4.4) is not ready until a volume is mounted and other conditions are met (see 4.3). While a direct access block device is not ready the device server shall

terminate logical block access commands with CHECK CONDITION status with the sense key set to NOT READY and the appropriate additional sense code for the condition.

Some direct access block devices may be switched from being ready to being not ready by using the START STOP UNIT command (see 5.25).

To make a direct access block device ready, an application client may be required to issue a START STOP UNIT command:

- a) with a START bit set to one and the POWER CONDITION field set to 0h (i.e., START_VALID); or
- b) with the POWER CONDITION field set to 1h (i.e., ACTIVE).

4.10 Initialization

A direct access block device may require initialization of its medium prior to processing logical block access commands. This initialization is requested by an application client using a FORMAT UNIT command (see 5.3). Parameters related to the format (e.g., logical block length) may be set with a MODE SELECT command (see SPC-4 and 6.5.2) prior to the format operation. Some direct access block devices are initialized by means outside the scope of this standard. The time when the initialization occurs is vendor specific.

Direct access block devices using a non-volatile medium may save the parameters related to the format and only require initialization once. However, some mode parameters may require initialization after each logical unit reset. A catastrophic failure of the direct access block device may require that an application client send a FORMAT UNIT command to recover from the failure.

Direct access block devices that use a volatile medium may require initialization after each logical unit reset prior to the processing of logical block access commands (see 4.2.2). Mode parameters may also require initialization after logical unit resets.

NOTE 2 - It is possible that mode parameter block descriptors read with a MODE SENSE command before a FORMAT UNIT completes contain information not reflecting the true state of the medium.

A direct access block device may become format corrupt after processing a MODE SELECT command that changes parameters (e.g., the logical block length) related to the medium format. During this time, the device server may terminate logical block access commands with CHECK CONDITION status with the sense key set to NOT READY and the appropriate additional sense code for the condition.

Any time the READ CAPACITY (10) parameter data (see 5.15.2) or the READ CAPACITY (16) parameter data (see 5.16.2) changes (e.g., when a FORMAT UNIT command or a MODE SELECT command causes a change to the logical block length or protection information, or when a vendor specific mechanism causes a change), then the device server shall establish a unit attention condition for the SCSI initiator port (see SAM-5) associated with each I_T nexus, except the I_T nexus on which the command causing the change was received with the additional sense code set to CAPACITY DATA HAS CHANGED.

NOTE 3 - Logical units compliant with SBC were not required to establish a unit attention condition with the additional sense code set to CAPACITY DATA HAS CHANGED.

4.11 Sanitize operations

4.11.1 Sanitize operations overview

A sanitize operation causes the device server to:

- a) affect the information on the logical unit's medium such that recovery of logical block data from the medium is not possible;
- b) affect the information in the cache by a method that is outside the scope of this standard such that previously existing data in cache is unable to be accessed; and

- c) prevent future access by an application client to cache or medium where the device server is unable to alter the information.

One of the following sanitize operations may be requested on the logical unit's medium:

- a) a sanitize overwrite operation that causes the device server to alter information by writing a data pattern to the medium one or more times;
- b) a sanitize block erase operation that causes the device server to alter information by setting the physical blocks to a vendor specific value; or
- c) a sanitize cryptographic erase operation that causes the device server to change encryption keys to prevent correct decryption of previously stored information, which may cause protection information, if any, to be indeterminate.

An application client may request that a sanitize operation be performed in the restricted completion mode or the unrestricted completion mode (see 4.11.3) using the AUSE bit in the SANITIZE command (see 5.24).

In the unrestricted completion mode, a SANITIZE command with the EXIT FAILURE MODE service action exits a failed sanitize operation.

In the restricted completion mode, the only method to exit a failed sanitize operation is for a SANITIZE command to request another sanitize operation and for that operation to complete without error. If a sanitize operation in the restricted completion mode completes with an error, and a subsequent SANITIZE command requests the unrestricted completion mode (i.e., the AUSE bit set to one), then the device server shall terminate that SANITIZE command as described in 5.24.1.

All sanitize operations shall be performed on:

- a) the medium that is being used to store logical block data;
- b) the medium that is not being used to store logical block data (e.g., areas previously used to store logical block data, areas available for allocation, and physical blocks that have become inaccessible); and
- c) all cache.

An application client requests that the device server perform a sanitize operation using the SANITIZE command (see 5.24). While the medium is write protected (see 4.12) the device server shall terminate a SANITIZE command with CHECK CONDITION status with the sense key set to DATA PROTECT and the appropriate additional sense code for the condition.

4.11.2 Performing a sanitize operation

Before performing a sanitize operation, the device server shall:

- a) terminate all commands in all task sets except INQUIRY commands, REPORT LUNS commands, LOG SENSE commands that specify the Temperature log page (see SPC-4), and REQUEST SENSE commands with CHECK CONDITION status with the sense key set to NOT READY, the additional sense code set to LOGICAL UNIT NOT READY, SANITIZE IN PROGRESS, and the PROGRESS INDICATION field in the sense data set to indicate that the sanitize operation is beginning;
- b) stop all enabled power condition timers (see SPC-4);
- c) stop all timers for enabled background scan operations (see 4.24);
- d) stop all timers or counters enabled for device specific background functions (see SPC-4); and
- e) discard partially downloaded microcode, if any.

While performing a sanitize operation, the device server shall:

- a) process INQUIRY commands, REPORT LUNS commands, LOG SENSE commands that specify the Temperature log page (see SPC-4), and REQUEST SENSE commands and terminate all other commands with CHECK CONDITION status with the sense key set to NOT READY, the additional sense code set to LOGICAL UNIT NOT READY, SANITIZE IN PROGRESS, and the PROGRESS INDICATION field in the sense data set to indicate the progress of the sanitize operation;
- b) provide pollable sense data (see SPC-4) with the sense key set to NOT READY, the additional sense code set to LOGICAL UNIT NOT READY, SANITIZE IN PROGRESS, and the PROGRESS INDICATION field set to indicate the progress of the sanitize operation;
- c) suspend the sanitize operation while processing the following conditions (see SAM-5):

- A) a power on;
- B) a hard reset;
- C) a logical unit reset; or
- D) a power loss expected;
- d) not suspend the sanitize operation while processing an I_T nexus loss;
- e) resume performing the sanitize operation after processing:
 - A) a logical unit reset; or
 - B) a power loss expected condition in which no power loss occurs within constraints defined by the applicable SCSI transport protocol standard (e.g., power loss timeout in SPL-3);
- f) process task management functions without affecting the processing of the sanitize operation (e.g., an ABORT TASK task management function aborts the SANITIZE command and has no effect on performing the sanitize operation);
- g) not alter mode data, INQUIRY data, or READ CAPACITY (16) parameter data (e.g., the number of logical blocks, logical block length, or protection information settings for the logical unit); and
- h) identify inaccessible physical blocks and in a vendor specific manner prevent future access to these blocks following a successful sanitize operation.

4.11.3 Completing a sanitize operation

If a sanitize operation completes without error, and logical block provisioning management (see 4.7.3) is supported, then:

- a) the initial condition for every LBA should be anchored (see 4.7.3.2) or deallocated (see 4.7.3.3); and
- b) read operations and write operations should complete without error.

If a sanitize operation completes without error and logical block provisioning management is not supported, then:

- a) read commands are processed as described in 5.24.2.2, 5.24.2.3, 5.24.2.4, and 5.24.2.5; and
- b) write operations should complete without error.

If the sanitize operation completes with an error in restricted completion mode, then the device server shall:

- a) terminate the SANITIZE command being performed, if any (e.g., the IMMED bit was set to zero in the CDB, and the failure occurs before status is returned for the command), with CHECK CONDITION status with the sense key set to MEDIUM ERROR and the additional sense code set to SANITIZE COMMAND FAILED; and
- b) until completion of a successful sanitize operation has occurred, terminate all commands, except SANITIZE commands allowed in the restricted completion mode, INQUIRY commands, REPORT LUNS commands, LOG SENSE commands specifying the Temperature log page (see SPC-4), and REQUEST SENSE commands, with CHECK CONDITION status with the sense key set to MEDIUM ERROR and the additional sense code set to SANITIZE COMMAND FAILED.

If a sanitize operation completes with an error in unrestricted completion mode, then the device server shall:

- a) terminate the SANITIZE command being performed, if any (e.g., the IMMED bit was set to zero in the CDB, and the failure occurs before status is returned for the command), with CHECK CONDITION status with the sense key set to MEDIUM ERROR and the additional sense code set to SANITIZE COMMAND FAILED; and
- b) until completion of a successful sanitize operation has occurred, terminate all commands, except SANITIZE commands, INQUIRY commands, REPORT LUNS commands, LOG SENSE commands specifying the Temperature log page (see SPC-4), and REQUEST SENSE commands, with CHECK CONDITION status with the sense key set to MEDIUM ERROR and the additional sense code set to SANITIZE COMMAND FAILED.

A sanitize operation that completed with error and was cleared with a SANITIZE command with the service action of EXIT FAILURE MODE may have not performed a complete sanitize operation (e.g., this action may enable the recovery of logical block data from the cache and medium for those logical blocks that were not sanitized).

After the sanitize operation completes the device server shall:

- 1) initialize all enabled timers and counters; and

- 2) start all enabled timers and counters for power conditions and background functions.

4.12 Write protection

Write protection prevents the alteration of the medium by logical block access commands issued to the device server. Write protection is controlled by:

- a) the user of the medium through manual intervention (e.g., a mechanical lock on the SCSI target device);
- b) hardware controls (e.g., tabs on the medium's housing); or
- c) software write protection.

All sources of write protection are independent. If present, any write protection shall cause otherwise valid logical block access commands that request alteration of the medium to be terminated by the device server with CHECK CONDITION status with the sense key set to DATA PROTECT and the appropriate additional sense code for the condition. Only when all write protections are disabled shall the device server process logical block access commands that request alteration of the medium.

Hardware write protection results when a physical attribute of the SCSI target device or its medium is changed to specify that writing shall be prohibited. Changing the state of the hardware write protection requires physical intervention, either with the SCSI target device or its medium. If allowed by the SCSI target device, then changing the hardware write protection while the medium is mounted results in vendor specific behavior that may include the writing of previously buffered data (e.g., data in cache).

Software write protection results when the device server is marked as write protected by the application client using the SWP bit in the Control mode page (see SPC-4). Changing the state of software write protection shall not prevent previously accepted logical block data (e.g., logical block data in cache) from being written to the medium.

The device server reports the status of write protection in the device server and on the medium with the DEVICE-SPECIFIC PARAMETER field in the mode parameter header (see 6.5.1).

4.13 Medium defects

4.13.1 Medium defects overview

Any medium has the potential for medium defects that cause data to be lost. Therefore, physical blocks and/or logical blocks may contain additional information that allows the detection of changes to the logical block data caused by medium defect or other phenomena. The additional information may also allow the logical block data to be reconstructed following the detection of such a change (e.g., ECC bytes).

A medium defect causes:

- a) a recovered error if the device server is able to read or write a logical block within the logical unit's recovery limits; or
- b) an unrecovered error if the device server is unable to read or write a logical block within the logical unit's recovery limits,

where the logical unit's recovery limits are:

- a) specified in the Read-Write Error Recovery mode page (see 6.5.8);
- b) specified in the Verify Error Recovery mode page (see 6.5.9); or
- c) vendor specific, if the device server does not implement the Read-Write Error Recovery mode page or the Verify Error Recovery mode page.

Direct access block devices may allow an application client to examine and modify the additional information by using the READ LONG commands (see 5.19 and 5.20) and the WRITE LONG commands (see 5.41 and 5.42). An application client may use the WRITE LONG commands to alter the additional information to test the

defect detection logic of the direct access block device or to emulate a logical block with an unrecovered read error when generating a mirror copy. This may induce a recovered error or an unrecovered error.

Direct access block devices may allow an application client to use the features of the WRITE LONG commands (see 5.41 and 5.42) to:

- a) disable error correction on specific logical blocks or physical blocks;
- b) disable automatic reassignment on specific logical blocks or physical blocks; and
- c) mark specific logical blocks or physical blocks as containing pseudo unrecovered errors with correction enabled or pseudo unrecovered errors with correction disabled.

These features provide methods for an application client to prevent recovered errors and unrecovered errors from being reported as informational exception conditions and prevent unnecessary reassign operations.

During a self-test operation (see SPC-4) or a background scan operation (see 4.24.1), the device server shall:

- a) ignore pseudo unrecovered errors with correction disabled; and
- b) process pseudo unrecovered errors with correction enabled.

The device server maintains the defect lists defined in table 9.

Table 9 – Defect lists (i.e., PLIST and GLIST)

Defect list	Source	Content
PLIST (i.e., primary defect list)	Manufacturer	Address descriptors (see 6.2) for physical blocks that contain permanent medium defects and never contain logical block data
GLIST (i.e., grown defect list)	FORMAT UNIT commands (see 5.3)	Address descriptors for physical blocks detected by the device server to have medium defects during an optional certification process performed during a format operation
		Address descriptors for physical blocks specified in the FORMAT UNIT parameter list (see 5.3.2)
	REASSIGN BLOCKS commands (see 5.21)	Address descriptors for physical blocks referenced by the LBAs specified in the reassign LBA list (see 5.21.2)
	Read medium operations	Address descriptors for physical blocks that have been reassigned as the result of automatic read reassignment
Write medium operations	Address descriptors for physical blocks that have been reassigned as the result of automatic write reassignment	

The READ DEFECT DATA commands (see 5.17 and 5.18) allow an application client to request that the device server return the PLIST and/or the GLIST.

The FORMAT UNIT command allows an application client to request that the device server clear the GLIST.

During a format operation, the device server shall not assign LBAs to any physical block in:

- a) the PLIST, if the PLIST is specified to be used; or
- b) the GLIST, if the GLIST is specified to be used.

A device server performs automatic reassignment of defects as specified by the settings in the Read-Write Error Recovery mode page (see 6.5.8).

The device server does not perform automatic read reassignment for an LBA referencing a logical block on which an unrecovered error has occurred. If the application client is notified by the device server that an unrecovered error occurred (e.g., as indicated by a read command being terminated with CHECK

CONDITION status with the sense key set to MEDIUM ERROR and the additional sense code set to UNRECOVERED READ ERROR) and:

- a) the application client is able to regenerate the logical block data for the LBA (e.g., in a redundancy group (see 4.19.1), the application client regenerates logical block data from the logical block data on the other logical units in the redundancy group) and the AWRE bit is set to one in the Read-Write Error Recovery mode page, then the application client may send a write command with that regenerated logical block data to trigger automatic write reassignment;
 - b) the application client is able to regenerate the logical block data for the LBA and the AWRE bit is set to zero in the Read-Write Error Recovery mode page, then the application client may:
 - 1) send a REASSIGN BLOCKS command to perform a reassign operation on the LBA; and
 - 2) send a write command with that regenerated logical block data;
- or
- c) the application client is unable to regenerate the logical block data for the LBA, then the application client may send a REASSIGN BLOCKS command to request that the device server perform a reassign operation on the LBA.

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC 14776-323:2017

4.13.2 Generation of defect lists

This standard defines address descriptor formats for describing defects (see 6.2). Table 10 lists the defects that each address descriptor format is capable of describing.

Table 10 – Address descriptor formats

Format	Single physical block	Multiple sequential physical blocks		Reference
		Entire track	Range	
Short block format	yes	no	no	6.2.2
Extended bytes from index format	yes ^a	yes ^e	yes ⁱ	6.2.3
Extended physical sector format	yes ^b	yes ^f	yes ⁱ	6.2.4
Long block format	yes	no	no	6.2.5
Bytes from index format	yes ^c	yes ^g	no	6.2.6
Physical sector format	yes ^d	yes ^h	no	6.2.7

^a Describes a single physical block with the MADS bit set to zero and the BYTES FROM INDEX field set to a value other than FFF_FFFFh.
^b Describes a single physical block with the MADS bit set to zero and the SECTOR NUMBER field set to a value other than FFF_FFFFh.
^c Describes a single physical block with the BYTES FROM INDEX field set to a value other than FFFF_FFFFh.
^d Describes a single physical block with the SECTOR NUMBER field set to a value other than FFFF_FFFFh.
^e Describes an entire track with the BYTES FROM INDEX field set to FFF_FFFFh.
^f Describes an entire track with the SECTOR NUMBER field set to FFF_FFFFh.
^g Describes an entire track with the BYTES FROM INDEX field set to FFFF_FFFFh.
^h Describes an entire track with the SECTOR NUMBER field set to FFFF_FFFFh.
ⁱ Describes a range with a pair of address descriptors using the same address descriptor format in which:

- the first address descriptor describes the starting location and has the MADS bit set to one;
- the second address descriptor describes the ending location and has the MADS bit set to zero; and
- the ending location is after the starting location.

For a direct access block device using rotating media (see 4.3.2), to represent two or more sequential physical blocks on the same track using a pair of address descriptors:

- the MADS bit shall be set to one in the first address descriptor;
- the MADS bit shall be set to zero in the second address descriptor;
- the CYLINDER NUMBER field in the first address descriptor shall be equal to the CYLINDER NUMBER field in the second address descriptor;
- the HEAD NUMBER field in the first address descriptor shall be equal to the HEAD NUMBER field in the second address descriptor;
- for a pair of extended bytes from index format address descriptors, the BYTES FROM INDEX field in the first address descriptor shall be less than the BYTES FROM INDEX field in the second address descriptor; and
- for a pair of extended physical sector format address descriptors, the SECTOR NUMBER field in the first address descriptor shall be less than the SECTOR NUMBER field in the second address descriptor.

For a direct access block device using rotating media, to represent two or more sequential tracks on the same head using a pair of address descriptors:

- the MADS bit shall be set to one in the first address descriptor;
- the MADS bit shall be set to zero in the second address descriptor;

- c) the CYLINDER NUMBER field in the first address descriptor shall be less than the CYLINDER NUMBER field in the second address descriptor;
- d) the HEAD NUMBER field in the first address descriptor shall be equal to the HEAD NUMBER field in the second address descriptor;
- e) for a pair of extended bytes from index format address descriptors, the BYTES FROM INDEX field in the first address descriptor and the second address descriptor shall be equal to FFF_FFFFh; and
- f) for a pair of extended physical sector format address descriptors, the SECTOR NUMBER field in the first address descriptor and the second address descriptor shall be equal to FFF_FFFFh.

4.14 Write and unmap failures

If one or more write commands have not completed when a power loss, a medium error, or hardware error occurs (e.g., a removable medium was incorrectly demounted), then any data in the logical blocks referenced by the LBAs specified by any of those commands is indeterminate. Before sending a read command or verify command specifying any LBAs that were specified by one of the write commands that did not complete, the application client should resend that write command. If an application client sends a read command or verify command specifying any LBAs that were specified by one of the write commands that did not complete before resending that write command, then the device server may return old data, new data, vendor specific data, or any combination thereof for the logical blocks referenced by the specified LBAs.

If logical block provisioning (see 4.7) is supported and one or more unmap commands have not completed when a power loss, medium error, or hardware error occurs, then the logical block provisioning state of the LBAs requested to be unmapped by any of those commands may or may not have changed. Before sending a read command or verify command specifying any LBAs that were specified by one of the unmap commands that did not complete, the application client should resend that unmap command.

4.15 Caches

4.15.1 Caches overview

Direct access block devices may implement caches. A cache is an area of temporary storage in the direct access block device (e.g., to enhance performance) separate from the medium that is not directly accessible by the application client.

A cache stores logical block data.

A cache may be volatile or non-volatile. A volatile cache does not retain data through power cycles. A non-volatile cache retains data through power cycles. There may be a limit on the amount of time a non-volatile cache is able to retain data without power (see 4.15.9).

4.15.2 Read caching

While processing read commands and verify commands, the device server may use the cache to store logical blocks that the application client may request at some future time. The algorithm used to manage the cache is not part of this standard. However, parameters are provided (see 6.5.5) to advise the device server about future requests, or to restrict the use of cache for a particular request.

4.15.3 Write caching

While processing write commands, the device server may perform a write cache operation to store logical block data that is to be written to the medium at a later time with a write medium operation. This is called write-back caching. A write command may complete prior to logical blocks being written to the medium. As a result of using write-back caching there is a period of time during which the logical block data may be lost if:

- a) power to the SCSI target device is lost and a volatile cache is being used; or
- b) a hardware failure occurs.

There is also the possibility of an error occurring during the subsequent write medium operation. If an error occurs during the write medium operation, then the error may be reported as a deferred error on a later command. The application client may request that write-back caching be disabled with the Caching mode page (see 6.5.5) to prevent detected write errors from being reported as deferred errors. Even with write-back caching disabled, undetected write errors may occur. Verify commands (e.g., VERIFY and WRITE AND VERIFY) may be used to detect those errors.

If processing a write command results in logical block data in cache that is different from the logical block data on the medium, then the device server shall retain that logical block data in cache until a write medium operation is performed using that logical block data. After the write medium operation is complete, the device server may retain that logical block data in cache.

4.15.4 Command interactions with caches

The application client may affect behavior of the cache with:

- a) the PRE-FETCH commands (see 5.8 and 5.9);
- b) the SYNCHRONIZE CACHE commands (see 5.26 and 5.27); and
- c) the Caching mode page (see 6.5.5).

When the cache becomes full of logical block data, the device server may replace the logical block data in the cache with new logical block data. The disable page out (DPO) bit in the CDBs of read commands, verify commands, and write commands allows the application client to influence the replacement of logical block data in the cache. A read command, verify command, or a write command with a DPO bit set to one is a hint to the device server that the logical blocks specified by that command are not likely to be accessed in the near future and should not be put in the cache or retained by the cache.

Application clients may use the force unit access (FUA) bit in the CDBs of read commands (e.g., see 5.11) or write commands (e.g., see 5.33) to specify that the device server shall access the medium or non-volatile cache.

Setting the DPO bit to one (e.g., see 5.11) and the FUA bit to one in all read commands and all write commands has the same effect as bypassing the volatile cache.

4.15.5 Write operation and write medium operation interactions with caches

For each LBA accessed by a write operation:

- 1) if a cache contains more recent logical block data for that LBA than the medium, then the device server shall:
 - A) perform a write cache operation to that LBA to update the logical block data in the cache;
 - B) invalidate that LBA in the cache and perform a write medium operation to that LBA; or
 - C) perform a write cache operation to that LBA to update the logical block data in the cache and perform a write medium operation to that LBA.

For each LBA accessed by a write medium operation that is not part of a write operation:

- 1) if a cache contains more recent logical block data for that LBA than the medium, then the device server shall:
 - A) perform a write cache operation to that LBA to update the logical block data in the cache; or
 - B) invalidate that LBA in the cache, before the device server performs the write medium operation to that LBA.

4.15.6 Read operation and read medium operation interactions with caches

For each LBA accessed by a read operation:

- 1) if a cache contains more recent logical block data for that LBA than the medium, then the device server shall perform a read cache operation from that LBA; or
- 2) the device server shall perform a read medium operation from that LBA.

For each LBA accessed by a read medium operation that is not part of a read operation:

- 1) if a cache contains more recent logical block data for the LBA than the medium, then the device server shall perform a write medium operation to that LBA; and
- 2) the device server may invalidate that LBA in the cache, before the device server performs the read medium operation from that LBA.

4.15.7 Verify medium operation interactions with caches

For each LBA accessed by a verify medium operation:

- 1) if a cache contains more recent logical block data for the LBA than the medium, then the device server shall perform a write medium operation to that LBA;
- 2) the device server may invalidate that LBA in the cache; and
- 3) before the device server performs the verify medium operation from that LBA.

4.15.8 Unmap operation interactions with caches

During an unmap operation, the device server changes any logical block data in the cache for the LBA unmapped by the operation so that any logical block data transferred by the device server to the Data-In Buffer while processing a subsequent read command reflects the results of the unmap operation (see 4.7.4.7.1 and 4.7.4.8.1).

4.15.9 Power loss effects on caches

The power, if any, needed to maintain a non-volatile cache may decrease to the point that the device server is unable to ensure the non-volatility of the cache for a vendor specific interval of time (e.g., the battery voltage becomes too low to sustain cache contents beyond a vendor specific time). If this occurs and the Extended INQUIRY Data VPD page (see SPC-4) indicates that the device server contains non-volatile cache (i.e., NV_SUP bit set to one), then:

- a) if the reporting of informational exceptions control warnings is enabled (i.e., the EWASC bit is set to one in the Information Exceptions Control mode page (see 6.5.6)), then the device server shall report the degraded non-volatile cache as specified in the Information Exceptions Control mode page with an additional sense code set to WARNING - DEGRADED POWER TO NON-VOLATILE CACHE; or
- b) if the reporting of informational exceptions control warnings is disabled (i.e., the EWASC bit is set to zero in the Information Exceptions Control mode page), then the device server shall establish a unit attention condition (see SAM-5) for the SCSI initiator port associated with every I_T nexus with the additional sense code set to WARNING - DEGRADED POWER TO NON-VOLATILE CACHE.

Non-volatile caches may become volatile (e.g., battery voltage becomes too low to sustain cache contents when power is lost). If non-volatile caches become volatile, then logical block data transferred for read commands or write commands in which the force unit access (FUA) bit in the CDB is set to one may bypass the cache.

If a non-volatile cache becomes volatile, then the device server shall set the REMAINING NON-VOLATILE TIME field to zero in the Non-volatile Cache log page (see 6.4.5).

If non-volatile cache becomes volatile and the Extended INQUIRY Data VPD page (see SPC-4) indicates that the device server contains non-volatile cache (i.e., the NV_SUP bit is set to one), then:

- a) if the reporting of informational exceptions control warnings is enabled (i.e., the EWASC bit is set to one in the Information Exceptions Control mode page (see 6.5.6)), then the device server shall report the change in the cache as specified in the Information Exceptions Control mode page with the additional sense code set to WARNING - NON-VOLATILE CACHE NOW VOLATILE; or
- b) if the reporting of informational exceptions control warnings is disabled (i.e., the EWASC bit is set to zero in the Information Exceptions Control mode page), then the device server shall establish a unit attention condition (see SAM-5) for the SCSI initiator port associated with every I_T nexus with the additional sense code set to WARNING - NON-VOLATILE CACHE NOW VOLATILE.

If:

- a) a power on or hard reset occurs;
- b) the Extended INQUIRY Data VPD page indicates that the device server contains a non-volatile cache (i.e., the NV_SUP bit is set to one); and
- c) the non-volatile cache is currently volatile,

then the device server shall establish a unit attention condition for the SCSI initiator port associated with every I_T nexus with the additional sense code set to WARNING - NON-VOLATILE CACHE NOW VOLATILE.

4.16 Implicit head of queue command processing

Each of the following commands defined by this standard may be processed by the task manager as if it has a HEAD OF QUEUE task attribute (see SAM-5), even if the command is received with a SIMPLE task attribute or an ORDERED task attribute:

- a) the READ CAPACITY (10) command (see 5.15); and
- b) the READ CAPACITY (16) command (see 5.16).

The following command defined by this standard shall be processed by the task manager as if it has a HEAD OF QUEUE task attribute, even if the command is received with a SIMPLE task attribute or an ORDERED task attribute:

- a) the SANITIZE command (see 5.24).

See SPC-4 for additional commands subject to implicit HEAD OF QUEUE command processing. See SAM-5 for additional rules on implicit head of queue processing.

4.17 Reservations

Reservation restrictions are placed on commands as a result of access qualifiers associated with the type of reservation. See SPC-4 for a description of reservations. The details of commands that are allowed under what types of reservations are described in table 11.

Commands from I_T nexuses holding a reservation should complete normally. Table 11 specifies the behavior of commands from registered I_T nexuses when a registrants only or all registrants type persistent reservation is present.

For each command, this standard or SPC-4 defines the conditions that result in the device server completing the command with RESERVATION CONFLICT status.

Table 11 – SBC-3 commands that are allowed in the presence of various reservations (part 1 of 2)

Command	Addressed logical unit has this type of persistent reservation held by another I_T nexus				
	From any I_T nexus		From registered I_T nexus (RR all types)	From I_T nexus not registered	
	Write Exclusive	Exclusive Access		Write Exclusive - RR	Exclusive Access - RR
COMPARE AND WRITE	Conflict	Conflict	Allowed	Conflict	Conflict
FORMAT UNIT	Conflict	Conflict	Allowed	Conflict	Conflict
GET LBA STATUS	Allowed	Conflict	Allowed	Allowed	Conflict
ORWRITE (16) / (32)	Conflict	Conflict	Allowed	Conflict	Conflict
POPULATE TOKEN	Allowed	Conflict	Allowed	Allowed	Conflict
PRE-FETCH (10) / (16)	Allowed	Conflict	Allowed	Allowed	Conflict
PREVENT ALLOW MEDIUM REMOVAL (Prevent=0)	Allowed	Allowed	Allowed	Allowed	Allowed
PREVENT ALLOW MEDIUM REMOVAL (Prevent ≠ 0)	Conflict	Conflict	Allowed	Conflict	Conflict
READ (10) / (12) / (16) / (32)	Allowed	Conflict	Allowed	Allowed	Conflict
READ CAPACITY (10) / (16)	Allowed	Allowed	Allowed	Allowed	Allowed
READ DEFECT DATA (10) / (12)	Allowed ^a	Conflict	Allowed	Allowed ^a	Conflict
READ LONG (10) / (16)	Conflict	Conflict	Allowed	Conflict	Conflict
REASSIGN BLOCKS	Conflict	Conflict	Allowed	Conflict	Conflict
REPORT REFERRALS	Allowed	Allowed	Allowed	Allowed	Allowed
SANITIZE	Conflict	Conflict	Allowed	Conflict	Conflict
START STOP UNIT with START bit set to one and POWER CONDITION field set to 0h	Allowed	Allowed	Allowed	Allowed	Allowed
START STOP UNIT with START bit set to zero or POWER CONDITION field set to a value other than 0h	Conflict	Conflict	Allowed	Conflict	Conflict
SYNCHRONIZE CACHE (10) / (16)	Conflict	Conflict	Allowed	Conflict	Conflict

Key:
RR = Registrants Only or All Registrants
Allowed = Commands received from I_T nexuses not holding the reservation or from I_T nexuses not registered when a registrants only or all registrants type persistent reservation is present should complete normally.
Conflict = Commands received from I_T nexuses not holding the reservation or from I_T nexuses not registered when a registrants only or all registrants type persistent reservation is present shall not be performed, and the device server shall complete the command with RESERVATION CONFLICT status.

^a The device server in logical units claiming compliance with SBC-2 may complete the command with RESERVATION CONFLICT status. Device servers may report whether certain commands are allowed in the PERSISTENT RESERVE IN command REPORT CAPABILITIES service action parameter data ALLOW COMMANDS field (see SPC-4).

Table 11 – SBC-3 commands that are allowed in the presence of various reservations (part 2 of 2)

Command	Addressed logical unit has this type of persistent reservation held by another I_T nexus				
	From any I_T nexus		From registered I_T nexus (RR all types)	From I_T nexus not registered	
	Write Exclusive	Exclusive Access		Write Exclusive - RR	Exclusive Access - RR
UNMAP	Conflict	Conflict	Allowed	Conflict	Conflict
VERIFY (10) / (12) / (16) / (32)	Allowed	Conflict	Allowed	Allowed	Conflict
WRITE (10) / (12) / (16) / (32)	Conflict	Conflict	Allowed	Conflict	Conflict
WRITE AND VERIFY (10) / (12) / (16) / (32)	Conflict	Conflict	Allowed	Conflict	Conflict
WRITE LONG (10) / (16)	Conflict	Conflict	Allowed	Conflict	Conflict
WRITE SAME (10) / (16) / (32)	Conflict	Conflict	Allowed	Conflict	Conflict
WRITE USING TOKEN	Conflict	Conflict	Allowed	Conflict	Conflict
XDWRITEREAD (10) / (32)	Conflict	Conflict	Allowed	Conflict	Conflict
XPWRITE (10) / (32)	Conflict	Conflict	Allowed	Conflict	Conflict
<p>Key: RR = Registrants Only or All Registrants Allowed = Commands received from I_T nexuses not holding the reservation or from I_T nexuses not registered when a registrants only or all registrants type persistent reservation is present should complete normally. Conflict = Commands received from I_T nexuses not holding the reservation or from I_T nexuses not registered when a registrants only or all registrants type persistent reservation is present shall not be performed, and the device server shall complete the command with RESERVATION CONFLICT status.</p>					
<p>^a The device server in logical units claiming compliance with SBC-2 may complete the command with RESERVATION CONFLICT status. Device servers may report whether certain commands are allowed in the PERSISTENT RESERVE IN command REPORT CAPABILITIES service action parameter data ALLOW COMMANDS field (see SPC-4).</p>					

4.18 Error reporting

4.18.1 Error reporting overview

If any of the conditions listed in table 12 occur during the processing of a command, then the device server shall terminate the command with CHECK CONDITION status with the sense key set to the specified value and the additional sense code set to the appropriate value for the condition. Some errors may occur after the completion status has already been reported. For such errors, SPC-4 defines a deferred error reporting mechanism. Table 12 lists some error conditions and the applicable sense keys. The list does not provide a complete list of all conditions that may cause CHECK CONDITION status.

Table 12 – Example error conditions

Condition	Sense key
Invalid LBA	ILLEGAL REQUEST
Unsupported option requested	ILLEGAL REQUEST
Logical unit reset, I_T nexus loss, or medium change since last command from this application client	UNIT ATTENTION
Logical block provisioning threshold notification	UNIT ATTENTION
Self diagnostic failed	HARDWARE ERROR
Unrecovered error	MEDIUM ERROR or HARDWARE ERROR
Recovered read error	RECOVERED ERROR
Pseudo unrecovered error	MEDIUM ERROR
Over-run or other error that may be resolved by repeating the command	ABORTED COMMAND
Attempt to write on write protected medium	DATA PROTECT

Direct access block devices compliant with this standard shall support both the fixed and descriptor formats of sense data (see SPC-4).

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC 14776-323:2017

Table 13 summarizes use of the sense data fields.

Table 13 – Sense data field usage for direct access block devices

Field	Usage	Reference
INFORMATION field ^a	READ LONG commands	5.19 and 5.20
	REASSIGN BLOCKS command	5.21
	WRITE LONG commands	5.41 and 5.42
	Any command that accesses the medium, based on the Read-Write Error Recovery mode page	4.20.3 and 6.5.8
	Any command that accesses the medium, based on the Verify Error Recovery mode page	
	Any command that is terminated with a logical block provisioning threshold notification	4.7.3.7.4
	COMPARE AND WRITE command	5.2
COMMAND-SPECIFIC INFORMATION field	EXTENDED COPY command	SPC-4
	REASSIGN BLOCKS command	5.21
	If rebuild assist mode is enabled (see 4.20), then any command that accesses the medium, based on the Read-Write Error Recovery mode page ^b	4.20.3
ILI bit	READ LONG commands	5.19 and 5.20
	WRITE LONG commands	5.41 and 5.42
^a See SPC-4 for a description of how the VALID bit interacts with the INFORMATION field. ^b If fixed format sense data is used but the value to be placed in the COMMAND-SPECIFIC INFORMATION field is greater than FFFF_FFFFh (e.g., an 8-byte LBA), then the device server shall report no value in the INFORMATION field (see SPC-4) and shall report no value in the COMMAND-SPECIFIC INFORMATION field (see SPC-4).		

If a command attempts to access or reference an invalid LBA, then the device server shall report the first invalid LBA (e.g., lowest numbered LBA) in the INFORMATION field of the sense data (see SPC-4).

If a recovered read error is reported, then device server shall report the last LBA (e.g., highest numbered LBA) on which a recovered read error occurred for the command in the INFORMATION field of the sense data (see SPC-4).

If an unrecovered error is reported, then the device server shall report the LBA of the logical block on which an unrecovered error occurred in the INFORMATION field of the sense data (see SPC-4).

4.18.2 Processing pseudo unrecovered errors

If a pseudo unrecovered error with correction disabled is encountered on a logical block (e.g., by a command, a background scan, or a background self-test), then the device server shall:

- perform no error recovery on the affected logical blocks, including any read error recovery enabled by the Read-Write Error Recovery mode page (see 6.5.8) or the Verify Error Recovery mode page (see 6.5.9);
- perform no automatic read reassignment or automatic write reassignment for the affected logical blocks, regardless of the settings of the AWRE bit and the ARRE bit in the Read-Write Error Recovery mode page;
- not consider errors on the affected logical blocks to be informational exception conditions as defined in the Information Exceptions Control mode page (see 6.5.6);

- d) not log errors on the affected logical blocks in any log page that contain error counters (see SPC-4); and
- e) in any information returned for the error (e.g., in sense data or in the Background Scan Results log page (see 6.4.2)), set the sense key to MEDIUM ERROR and either:
 - A) should set the additional sense code to READ ERROR – LBA MARKED BAD BY APPLICATION CLIENT; or
 - B) may set the additional sense code to UNRECOVERABLE READ ERROR.

The logical unit shall clear a pseudo unrecovered error if it processes or performs one of the following for that LBA:

- a) a format operation;
- b) a reassign operation;
- c) a sanitize overwrite operation;
- d) a sanitize block erase operation; or
- e) a write command that is not a WRITE LONG command specifying a pseudo unrecovered error.

The logical unit may clear a pseudo unrecovered error if it processes or performs one of the following for that LBA:

- a) a sanitize cryptographic erase operation;
- b) an unmap operation; or
- c) a MODE SELECT command that uses the mode parameter block descriptor (see 6.5.2) to change the capacity to be lower than that LBA.

The logical unit shall not clear a pseudo unrecovered error if it processes one of the following for that LBA:

- a) a read command.

If a pseudo unrecovered error with correction enabled is encountered on a logical block (e.g., by a command, a background scan, or a background self-test), then the device server shall:

- a) if enabled by the Read-Write Error Recovery mode page (see 6.5.8) or the Verify Error Recovery mode page (see 6.5.9), perform error recovery on the affected logical blocks;
- b) perform no automatic read reassignment or automatic write reassignment for the affected logical blocks, regardless of the settings of the AWRE bit and the ARRE bit in the Read-Write Error Recovery mode page;
- c) consider errors on the affected logical blocks to be informational exception conditions as defined in the Information Exceptions Control mode page (see 6.5.6);
- d) log errors on the affected logical blocks in any log page that contain error counters (see SPC-4); and
- e) in any information returned for the error (e.g., in sense data or in the Background Scan Results log page (see 6.4.2)), set the sense key to MEDIUM ERROR and either:
 - A) should set the additional sense code to READ ERROR – LBA MARKED BAD BY APPLICATION CLIENT; or
 - B) may set the additional sense code to UNRECOVERABLE READ ERROR.

4.18.3 Block commands sense data descriptor

Table 14 defines the block commands sense data descriptor used in descriptor format sense data (see SPC-4) for direct access block devices.

Table 14 – Block commands sense data descriptor format

Byte	Bit	7	6	5	4	3	2	1	0
0	DESCRIPTOR TYPE (05h)								
1	ADDITIONAL LENGTH (02h)								
2	Reserved								
3	Reserved			ILI		Reserved			

The DESCRIPTOR TYPE field and the ADDITIONAL LENGTH field are defined in SPC-4 and shall be set to the values shown in table 14 for the block commands sense data descriptor.

The incorrect length indication (ILI) bit set to one indicates that:

- the command was a READ LONG command or WRITE LONG command with a requested data length; and
- the requested data length in the READ LONG command or WRITE LONG command did not match the length of the logical block.

An ILI bit set to zero indicates that either:

- the command was not a READ LONG command or WRITE LONG command with a requested data length; or
- the requested data length in the READ LONG command or WRITE LONG command matched the length of the logical block.

4.18.4 User data segment referral sense data descriptor

Table 15 defines the user data segment referral sense data descriptor used in descriptor format sense data for direct access block devices. The user data segment referral sense data descriptor contains descriptors indicating the user data segment(s) on the logical unit and the SCSI target port groups through which those user data segments may be accessed (see 4.28).

Table 15 – User data segment referral sense data descriptor format

Byte	Bit	7	6	5	4	3	2	1	0
0	DESCRIPTOR TYPE (0Bh)								
1	ADDITIONAL LENGTH (y -1)								
2	Reserved								NOT_ALL_R
3	Reserved								
User data segment referral descriptor list									
4	User data segment referral descriptor [first]								
...									
4 + n									
⋮									
y - m	User data segment referral descriptor [last]								
...									
y									

The DESCRIPTOR TYPE field is defined in SPC-4 and shall be set to the value shown in table 15 for the user data segment referral sense data descriptor.

The ADDITIONAL LENGTH field indicates the number of bytes that follow in the logical block referrals sense data descriptor.

A not all referrals (NOT_ALL_R) bit set to zero indicates that the list of user data segment referral descriptors is a complete list of user data segments. A NOT_ALL_R bit set to one indicates that there are more user data segments than are able to be indicated by the user data segment referral sense data.

Each user data segment referral descriptor (see table 16) indicates information identifying:

- a) a user data segment that is accessible through the SCSI target port groups indicated by this descriptor; and
- b) one or more SCSI target port groups through which the user data segment indicated by this descriptor is able to be accessed.

User data segment referral descriptors shall be listed in ascending LBA order. If a user data segment referral descriptor describes the last user data segment (i.e., points to the largest LBA) and the preceding user data segment descriptors do not represent the complete list of user data segments, then the next user data segment referral descriptor, if any, shall describe the first user data segment (i.e., the user data segments may wrap).

Table 16 defines the user data segment referral descriptor.

Table 16 – User data segment referral descriptor format

Byte	Bit	7	6	5	4	3	2	1	0
0		Reserved							
...									
2									
3		NUMBER OF TARGET PORT GROUP DESCRIPTORS							
4	(MSB)	FIRST USER DATA SEGMENT LBA							
...									
11									
12	(MSB)	LAST USER DATA SEGMENT LBA							
...									
19									
Target port group descriptor list									
20		Target port group descriptor [first]							
...									
23									
⋮									
m-3		Target port group descriptor [last]							
...									
m									

The NUMBER OF TARGET PORT GROUP DESCRIPTORS field indicates the number of target port group descriptors that follow.

The FIRST USER DATA SEGMENT LBA field indicates the first LBA of the first user data segment (see 4.28) indicated by this descriptor.

The LAST USER DATA SEGMENT LBA field indicates the last LBA of the last user data segment (see 4.28) indicated by this descriptor.

The target port group descriptor (see table 17) specifies the target port group and the asymmetric access state of the target port group (see SPC-4). The device server shall return one target port group descriptor for each target port group in a target port asymmetric access state of active/optimized, active/non-optimized, or transitioning. The device server may return one target port group descriptor for each target port group in a target port asymmetric access state of unavailable.

Table 17 – Target port group descriptor

Byte	Bit	7	6	5	4	3	2	1	0
0		Reserved				ASYMMETRIC ACCESS STATE			
1		Reserved							
2	(MSB)	TARGET PORT GROUP							
3		(LSB)							

The ASYMMETRIC ACCESS STATE field (see SPC-4) contains the asymmetric access state of the user data segment(s) specified by this descriptor that may be accessed through this target port group.

The TARGET PORT GROUP field specifies a target port group (see SPC-4) that the application client uses when issuing commands associated with the user data segments specified by this descriptor.

4.18.5 Direct-access block device sense data descriptor

Table 18 defines the direct-access block device sense data descriptor, which may be used in descriptor format sense data (see SPC-4) instead of any of the following sense data descriptors:

- a) information (see SPC-4);
- b) command-specific information (see SPC-4);
- c) sense key specific (see SPC-4);
- d) field replaceable unit (see SPC-4); and
- e) block commands (see 4.18.3).

If the device server includes the direct-access block device sense data descriptor in a sense data descriptor list, then it shall not include any of those sense data descriptors in the same sense data descriptor list.

Table 18 – Direct-access block device sense data descriptor format

Byte	Bit	7	6	5	4	3	2	1	0
0	DESCRIPTOR TYPE (0Dh)								
1	ADDITIONAL LENGTH (1Eh)								
2	VALID	Reserved		ILI	Reserved				
3	Reserved								
4	SKSV	Sense key specific information							
...									
6									
7	FIELD REPLACEABLE UNIT CODE								
8	(MSB)	INFORMATION							
...									
15									
16	(MSB)	COMMAND-SPECIFIC INFORMATION							
...									
23									

The DESCRIPTOR TYPE field is described in SPC-4, and shall be set as shown in table 18 for the direct-access block device sense data descriptor.

A VALID bit set to zero indicates that the INFORMATION field is not defined in this standard or any other command standard. A VALID bit set to one indicates the INFORMATION field contains valid information as defined in this standard or a command standard.

The ILI bit is described in the block commands sense data descriptor (see 4.18.3).

A sense-key specific valid (SKSV) bit set to one indicates the sense key specific information contains valid information as defined in SPC-4. An SKSV bit set to zero indicates that the sense key specific information is not as defined by SPC-4.

The sense key specific information is described in the sense key specific sense data descriptor (see SPC-4).

The FIELD REPLACEABLE UNIT CODE field is described in the field replaceable unit sense data descriptor (see SPC-4).

The INFORMATION field is described in the information sense data descriptor (see SPC-4).

The COMMAND-SPECIFIC INFORMATION field is described in the command-specific information sense data descriptor (see SPC-4). The COMMAND-SPECIFIC INFORMATION field should be ignored in sense data for a command or operation for which the COMMAND-SPECIFIC INFORMATION field is not defined and in sense data that is not related to a command or operation (e.g., pollable sense data).

4.19 Model for XOR commands

4.19.1 Model for XOR commands overview

In storage arrays, a SCSI storage array device (see SCC-2) organizes a group of direct access block devices into objects. The type of object supported by this model is the redundancy group, where some of the logical blocks on the direct access block devices are used for XOR-protected data and some of the logical blocks are used for check data. The check data is generated by performing a cumulative XOR operation of the

XOR-protected data. The XOR operation may be performed by the SCSI storage array device or by the direct access block devices.

A direct access block device containing XOR-protected data is called a data disk. A direct access block device containing check data is called a parity disk.

Performing the XOR operation in the direct access block devices may result in a reduced number of data transfers across a service delivery subsystem.

For example, when the XOR operation is performed within the SCSI storage array device, the following commands are used for a typical update write sequence:

- 1) a read command to the data disk;
- 2) a write command to the data disk;
- 3) a read command to the parity disk; and
- 4) a write command to the parity disk.

The SCSI storage array device also does two internal XOR operations in this sequence.

In contrast, during SCSI storage array device supervised XOR operations (see 4.19.2) only the following operations are used:

- a) a write command to the data disk;
- b) a read command to the data disk; and
- c) a write command to the parity disk.

4.19.2 SCSI storage array device supervised XOR operations

4.19.2.1 SCSI storage array device supervised XOR operations overview

A SCSI storage array device (see SCC-2) supervises three basic operations that require XOR functionality:

- a) update write operation (see 4.19.2.2);
- b) regenerate operation (see 4.19.2.3); and
- c) rebuild operation (see 4.19.2.4).

Command sequences for each of these operations use the device servers in the direct access block devices to perform the necessary XOR functions.

XDWRITEREAD commands (see 5.47 and 5.48) and XPWRITE commands (see 5.49 and 5.50) are required to implement SCSI storage array device supervised XOR operations. The SCSI storage array device also uses READ commands and WRITE commands for certain operations.

4.19.2.2 Update write operation

The update write operation writes new XOR-protected data to a data disk and updates the check data on the parity disk. The sequence performed by the SCSI storage array device (see SCC-2) is as follows:

- 1) send an XDWRITEREAD command to the data disk with the Data-Out Buffer containing the new XOR-protected data; and
- 2) send an XPWRITE command to the parity disk with the Data-Out Buffer for this command containing the logical block data from the Data-In Buffer of the XDWRITEREAD command.

4.19.2.3 Regenerate operation

The regenerate operation is used to recreate one or more logical blocks that have an error. This is accomplished by reading the associated logical block from each of the other direct access block devices within the redundancy group and performing an XOR operation on each of these logical blocks. The last XOR result is the data that should have been present on the unreadable direct access block device. The number of steps is dependent on the number of direct access block devices in the redundancy group, but the sequence performed by the SCSI storage array device (see SCC-2) is as follows:

- 1) send a READ command to the first direct access block device;

- 2) send an XDWRITEREAD command with the DISABLE WRITE bit set to one to the next direct access block device using the logical block data from the Data-In Buffer of the previous command as the logical block data in the Data-Out Buffer for this command; and
- 3) repeat step 2) until all direct access block devices in the redundancy group except the failed device have been accessed. The logical block data in the Data-In Buffer for the last command is the regenerated data.

4.19.2.4 Rebuild operation

The rebuild operation is similar to the regenerate operation, except that the last XOR result is written to the replacement device. This function is used when a failed device is replaced and the SCSI storage array device (see SCC-2) is writing the rebuilt data to the replacement device. The number of steps is dependent on the number of direct access block devices in the redundancy group, but the sequence performed by the SCSI storage array device (see SCC-2) is as follows:

- 1) perform the regenerate operation (see 4.19.2.3); and
- 2) send a WRITE command to the replacement device using the regenerated data.

4.19.3 Array subsystem considerations

4.19.3.1 Array subsystem considerations overview

This subclause lists considerations that apply to any array subsystem and describes how use of the XOR commands may affect handling of those situations.

4.19.3.2 Access to an inconsistent stripe

A stripe is a set of corresponding extents from two or more direct access block devices.

When the SCSI storage array device (see SCC-2) issues an update write to a data disk, the data in the data disk has been updated when successful status is returned for the command. Until the parity disk has been updated the associated stripe in the redundancy group is not consistent (i.e., performing an XOR operation on the XOR-protected data does not produce the check data).

The SCSI storage array device shall keep track of this window of inconsistency and ensure that a regenerate or rebuild operation for any data extent within the stripe is not attempted until after the parity disk has been updated, making the stripe consistent again. For multi-initiator systems, tracking the updates may be more complex because each SCSI storage array device is required to ensure that a second SCSI storage array device is not writing to a stripe that the first SCSI storage array device is regenerating or rebuilding. The coordination between SCSI storage array devices is system specific and is beyond the scope of this standard.

If a device server terminates any of the XOR commands with CHECK CONDITION status and an unrecovered error is indicated, then an inconsistent stripe may result. It is the SCSI storage array device's responsibility to identify the failing device, identify the scope of the failure, and then to limit access to the inconsistent stripe. The recovery procedures that the SCSI storage array device implements are outside the scope of this standard.

4.20 Rebuild assist mode

4.20.1 Rebuild assist mode overview

The rebuild assist mode provides a method for a SCSI storage array device (see SCC-2) to read recovered logical blocks from a failed logical unit in a storage array instead of rebuilding the logical blocks from other logical units in the storage array. This mode allows the failed logical unit to report logical blocks that are unreadable without requiring the SCSI storage array device to read every LBA in the failed logical unit to determine the unrecovered logical blocks. The SCSI storage array device then copies the logical blocks recovered from the failed logical unit to a replacement logical unit and only rebuilds the failed logical blocks.

Enabling the rebuild assist mode:

- a) may cause the device server to initiate a self test to identify the scope of failures, if any;
- b) modifies READ command recovery behavior by the device server based on the setting of the RARC bit (see 4.20.3 and 5.11); and
- c) may cause sense data to be returned by the device server that indicates the location of multiple failing logical blocks on read commands and write commands.

The self-test operations performed by the device server while rebuild assist mode is enabled may result in detection of failed physical elements. A predicted unrecovered error is an unrecovered error that is the result of an attempt by the device server to access an LBA associated with a failed physical element. An unpredicted unrecovered error is an unrecovered error that is the result of a device server accessing an LBA that is not associated with a failed physical element.

4.20.2 Enabling rebuild assist mode

An application client should enable rebuild assist mode after the application client determines that a rebuild is required. The application client enables the rebuild assist mode by setting the ENABLED bit to one and setting the DISABLED PHYSICAL ELEMENT field to all zeros in the Rebuild Assist Output diagnostic page (see 6.3.3).

If a SEND DIAGNOSTIC command requests the enabling of the rebuild assist mode, then the device server:

- 1) shall enable the rebuild assist mode;
- 2) may perform a diagnostic test of the physical elements contained within the logical unit; and
- 3) should disable any physical elements that are not functional if a diagnostic test of the physical elements is performed.

The application client may verify that rebuild assist mode is enabled by verifying that the ENABLED bit is set to one in the Rebuild Assist Input diagnostic page (see 6.3.2).

4.20.3 Using the rebuild assist mode

4.20.3.1 Using rebuild assist mode overview

After rebuild assist mode is enabled, the application client should issue read commands to read the available logical block data from the failed logical unit. If the device server does not detect an unrecovered error while processing a read command, then the device server should complete the read command without error.

The rebuild assist mode allows the device server to report unrecovered read errors or unrecovered write errors that are either predicted (i.e., predicted unrecovered errors) or unpredicted (i.e., unpredicted unrecovered errors).

4.20.3.2 Unpredicted unrecovered read error

If a device server receives a read command with the RARC bit set to one, then rebuild assist mode shall not affect processing of the read command.

If rebuild assist mode is enabled and a device server receives a read command with the RARC bit set to zero and the device server detects an unpredicted unrecovered error that is not a pseudo unrecovered read error (see 4.18.2), then the device server:

- 1) shall perform limited read recovery that is vendor specific;
- 2) shall transfer the data for all recovered logical blocks, if any, from the logical block referenced by the starting LBA of the failed read command up to the first unrecovered logical block (i.e., the lowest numbered LBA) to the Data-in Buffer;
- 3) shall terminate the command with CHECK CONDITION status with the sense key set to MEDIUM ERROR, the additional sense code set to UNRECOVERED READ ERROR, and report the LBA referencing the unrecovered logical block in the INFORMATION field (see SPC-4); and
- 4) may use this failure in a vendor specific manner to predict other logical blocks that may be unrecovered.

If the application client receives sense data with the sense key set to MEDIUM ERROR, the additional sense code set to UNRECOVERED READ ERROR, and the INFORMATION field indicating a valid LBA (see SPC-4), then the application client should issue the next read command with the starting LBA set to the contents of the INFORMATION field plus one.

4.20.3.3 Predicted unrecovered read error

If the device server receives a read command with the RARC bit set to one, then rebuild assist mode shall not affect the processing of the read command.

If rebuild assist mode is enabled and the device server receives a read command with the RARC bit set to zero, and the device server detects a predicted unrecovered error, then the device server:

- 1) shall perform limited read recovery that is vendor specific;
- 2) shall transfer the data for all recovered logical blocks, if any, from the logical block referenced by the starting LBA of the failed read command up to the first predicted unrecovered LBA (i.e., the lowest numbered LBA) to the Data-in Buffer; and
- 3) shall terminate the read command with CHECK CONDITION status with the sense key set to ABORTED COMMAND, the additional sense code set to MULTIPLE READ ERRORS, and:
 - A) report the following value in the INFORMATION field (see SPC-4):
 - a) the LBA referencing the first unrecovered logical block (i.e., the lowest numbered LBA);and
 - B) report the following value in the COMMAND-SPECIFIC INFORMATION field (see SPC-4):
 - a) the LBA referencing the last unrecovered logical block (i.e., the highest numbered LBA) in a sequence of contiguous unrecovered logical blocks that started with the LBA indicated in the INFORMATION field.

If the application client receives sense data with the sense key set to ABORTED COMMAND, the additional sense code set to MULTIPLE READ ERRORS, and the INFORMATION field indicating a valid LBA (see SPC-4), then the application client should issue the next read command with the starting LBA set to the contents of the COMMAND-SPECIFIC INFORMATION field plus one to continue recovering data from the logical unit. This process should be repeated until all of the LBAs have been scanned.

4.20.3.4 Unpredicted unrecovered write error

If rebuild assist mode is enabled and the device server detects an unpredicted unrecovered error while processing a write command, then the device server shall terminate the command with CHECK CONDITION status with the sense key set to MEDIUM ERROR, the additional sense code set to WRITE ERROR, and report the LBA referencing the first logical block (i.e., the lowest numbered LBA) in error in the INFORMATION field (see SPC-4).

4.20.3.5 Predicted unrecovered write error

If rebuild assist mode is enabled and the device server detects a predicted unrecovered error while processing a write command, then the device server:

- 1) transfers the write data from the Data-Out Buffer;
- 2) writes the transferred data up to the logical block referenced by the failing LBA; and
- 3) shall terminate the write command with CHECK CONDITION status with the sense key set to ABORTED COMMAND, the additional sense code set to MULTIPLE WRITE ERRORS, and:
 - A) report the following value in the INFORMATION field (see SPC-4):
 - a) the LBA referencing the first logical block (i.e., the lowest numbered LBA) in error;and
 - B) report the following value in the COMMAND-SPECIFIC INFORMATION field (see SPC-4):
 - a) the LBA referencing the last logical block (i.e., the highest numbered LBA) in error in a sequence of contiguous logical blocks that started with the LBA indicated in the INFORMATION field.

If the application client receives sense data with the sense key set to ABORTED COMMAND, the additional sense code set to MULTIPLE WRITE ERRORS, and the INFORMATION field indicating a valid LBA (see SPC-4), then the application client should issue the next write command with the starting LBA set to the contents of the COMMAND-SPECIFIC INFORMATION field plus one to continue writing to the logical unit.

4.20.4 Disabling the rebuild assist mode

Rebuild assist mode shall be disabled after a power on.

Rebuild assist mode shall not be affected by a hard reset, an I_T nexus loss, or any task management functions (see SAM-5).

The application client disables rebuild assist mode by setting the ENABLED bit to zero in the Rebuild Assist Output diagnostic page (see 6.3.3).

4.20.5 Testing rebuild assist mode

The Rebuild Assist Output diagnostic page (see 6.3.3) provides a method to test the application client's rebuild process.

An application client places a logical unit into a simulated failing condition by setting the ENABLED bit to one and setting one or more bits in the DISABLED PHYSICAL ELEMENT field to one in the Rebuild Assist Output diagnostic page.

Each bit in the DISABLED PHYSICAL ELEMENT field represents a physical element that is associated with a group of LBAs that are treated as having predicted unrecovered read errors. The correlation of bits in the DISABLED PHYSICAL ELEMENT field to LBAs in the logical unit is vendor specific.

An application client ends a test by disabling the rebuild assist mode (see 4.20.4).

4.21 START STOP UNIT and power conditions

4.21.1 START STOP UNIT and power conditions overview

The START STOP UNIT command (see 5.25) allows an application client to control the power condition of a logical unit. This method includes specifying that the logical unit transition to a specific power condition.

In addition to the START STOP UNIT command, the power condition of a logical unit may be controlled by the Power Condition mode page (see SPC-4). If both the START STOP UNIT command and the Power Condition mode page methods are being used to control the power condition of the same logical unit, then the power condition specified by any START STOP UNIT command shall override the Power Condition mode page's power control.

4.21.2 Processing of concurrent START STOP UNIT commands

If a START STOP UNIT command is being processed by the device server, and a subsequent START STOP UNIT command for which the CDB is validated requests that the logical unit change to a different power condition than was specified by the START STOP UNIT command being processed, then the device server shall terminate the subsequent START STOP UNIT command with CHECK CONDITION status with the sense key set to NOT READY and the additional sense code set to LOGICAL UNIT NOT READY, START STOP UNIT COMMAND IN PROGRESS.

The constraints on concurrent START STOP UNIT commands apply only to commands that have the IMMED bit set to zero. The effects of concurrent power condition changes requested by START STOP UNIT commands with the IMMED bit set to one are vendor specific.

4.21.3 Managing logical block access commands during a change to the active power condition

Application clients may minimize the return of BUSY status or TASK SET FULL status during a change to the active power condition by:

- a) polling the power condition using the REQUEST SENSE command (see SPC-4); or
- b) sending a START STOP UNIT command with the IMMED bit set to zero and the START bit set to one and waiting for GOOD status to be returned.

4.21.4 Stopped power condition

In addition to the active power condition, idle power conditions, and standby power conditions described in SPC-4, this standard describes the stopped power condition.

While in the stopped power condition:

- a) the device server shall terminate TEST UNIT READY commands and logical block access commands with CHECK CONDITION status with the sense key set to NOT READY and the additional sense code set to LOGICAL UNIT NOT READY, INITIALIZING COMMAND REQUIRED;
- b) the power consumed by the SCSI target device while in the stopped power condition should be less than the power consumed when the logical unit is in the active power condition or any of the idle power conditions (e.g., for direct access block devices that have a rotating medium, the medium is stopped in the stopped power condition);
- c) the peak power consumption during a change from the stopped power condition to the active power condition or an idle power condition is not limited by this standard; and
- d) the peak power consumption during a change from the stopped power condition to a standby power condition shall be no more than the typical peak power consumption in the active power condition.

No power condition defined in this standard shall affect the supply of any power required for proper operation of a service delivery subsystem.

4.21.5 START STOP UNIT and power condition state machine

4.21.5.1 START STOP UNIT and power condition state machine overview

The SSU_PC (START STOP UNIT and power condition) state machine for logical units implementing the START STOP UNIT command describes the:

- a) logical unit power states and transitions resulting from specifications in the START STOP UNIT command;
- b) settings in the Power Condition mode page (see SPC-4); and
- c) the processing of commands.

NOTE 4 - The SSU_PC state machine is an enhanced version of the PC (power condition) state machine described in SPC-4.

The SSU_PC state machine consists of the states shown in table 19.

Table 19 – Summary of states in the SSU_PC state machine

State	Reference	PC state machine state with additional definition (see SPC-4)
SSU_PC0:Powered_On ^a	4.21.5.2	PC0:Powered_On
SSU_PC1:Active	4.21.5.3	PC1:Active
SSU_PC2:Idle	4.21.5.4	PC2:Idle
SSU_PC3:Standby	4.21.5.5	PC3:Standby
SSU_PC4:Active_Wait	4.21.5.6	PC4:Active_Wait
SSU_PC5:Wait_Idle	4.21.5.7	PC5:Wait_Idle
SSU_PC6:Wait_Standby	4.21.5.8	PC6:Wait_Standby
SSU_PC7:Idle_Wait	4.21.5.9	n/a
SSU_PC8:Stopped	4.21.5.10	n/a
SSU_PC9:Standby_Wait	4.21.5.11	n/a
SSU_PC10:Wait_Stopped	4.21.5.12	n/a
^a SSU_PC0:Powered_On is the initial state.		

While in the following SSU_PC states, the logical unit may be increasing power usage to enter a higher power condition:

- a) SSU_PC4:Active_Wait;
- b) SSU_PC7:Idle_Wait; and
- c) SSU_PC9:Standby_Wait.

While in the following SSU_PC states, the logical unit may be decreasing power usage to enter a lower power condition:

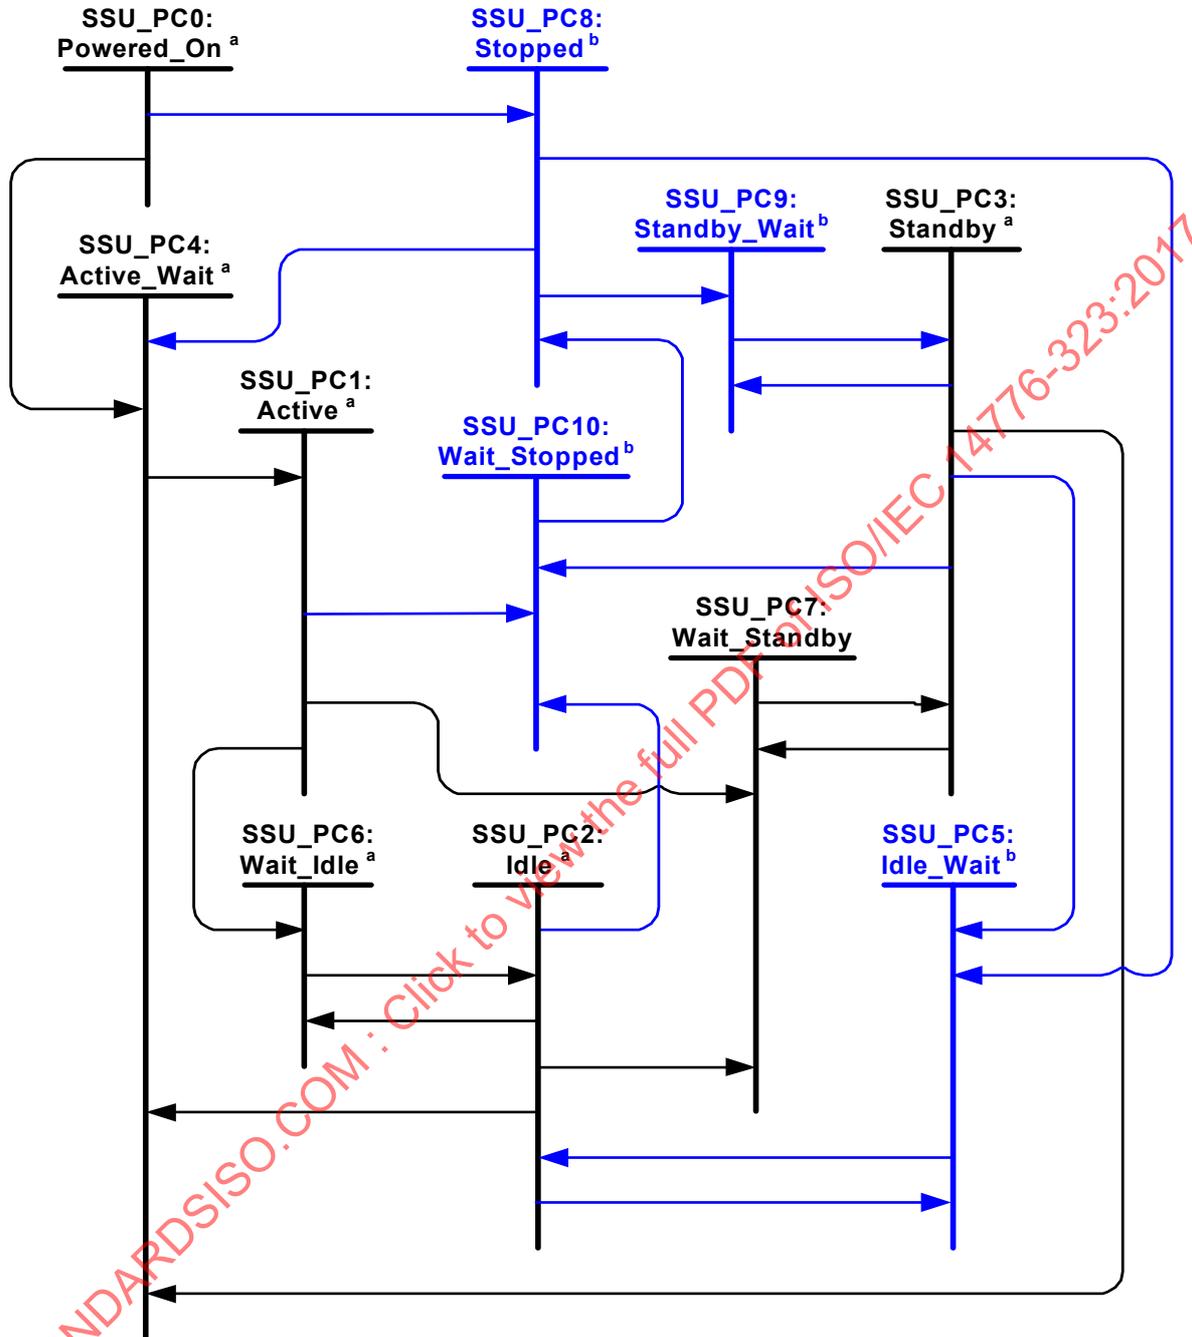
- a) SSU_PC5:Wait_Idle;
- b) SSU_PC6:Wait_Standby; and
- c) SSU_PC10:Wait_Stopped.

Any command causing a state machine transition (e.g., a START STOP UNIT command with the IMMED bit set to zero) shall not complete with GOOD status until this state machine reaches the state (i.e., power condition) required or specified by the command.

The SSU_PC state machine shall start in the SSU_PC0:Powered_On state after power on. The SSU_PC state machine shall be configured to transition to the SSU_PC4:Active_Wait state or the SSU_PC8:Stopped state after power on by a mechanism outside the scope of this standard.

This state machine references timers controlled by the Power Condition mode page (see SPC-4) and refers to the START STOP UNIT command (see 5.25).

Figure 12 describes the SSU_PC state machine.



^a This state or transition is also described in SPC-4, but may have additional characteristics unique to this standard (e.g., a transition to or from a state described in this standard).
^b This state or transition is described in this standard.

Figure 12 – SSU_PC state machine

4.21.5.2 SSU_PC0:Powered_On state

4.21.5.2.1 SSU_PC0:Powered_On state description

See the PC0:Powered_On state in SPC-4 for details about this state.

4.21.5.2.2 Transition SSU_PC0:Powered_On to SSU_PC4:Active_Wait

This transition shall occur if:

- a) the logical unit is ready to begin power on initialization; and
- b) the logical unit has been configured to transition to the SSU_PC4:Active_Wait state.

The transition shall include a Transitioning From Powered On argument.

4.21.5.2.3 Transition SSU_PC0:Powered_On to SSU_PC8:Stopped

This transition shall occur if:

- a) the logical unit has been configured to transition to the SSU_PC8:Stopped state.

The transition shall include a Transitioning From Powered On argument.

4.21.5.3 SSU_PC1:Active state**4.21.5.3.1 SSU_PC1:Active state description**

See the PC1:Active state in SPC-4 for details about this state.

4.21.5.3.2 Transition SSU_PC1:Active to SSU_PC5:Wait_Idle

This transition shall occur if:

- a) the device server processes a START STOP UNIT command (see 5.25) with the POWER CONDITION field set to 2h (i.e., IDLE); or
- b) an idle condition timer (see SPC-4) is enabled, and that timer has expired.

The transition shall include a:

- a) Transitioning To Idle_a argument, if:
 - A) the highest priority timer that expired is the idle_a condition timer; or
 - B) the START STOP UNIT command being processed has the POWER CONDITION MODIFIER field set to 0h (i.e., idle_a power condition);
 - b) Transitioning To Idle_b argument, if:
 - A) the highest priority timer that expired is the idle_b condition timer; or
 - B) the START STOP UNIT command being processed has the POWER CONDITION MODIFIER field set to 1h (i.e., idle_b power condition);
- or
- c) Transitioning To Idle_c argument, if:
 - A) the highest priority timer that expired is the idle_c condition timer; or
 - B) the START STOP UNIT command being processed has the POWER CONDITION MODIFIER field set to 2h (i.e., idle_c power condition).

4.21.5.3.3 Transition SSU_PC1:Active to SSU_PC6:Wait_Standby

This transition shall occur if:

- a) the device server processes a START STOP UNIT command (see 5.25) with the POWER CONDITION field set to 3h (i.e., STANDBY); or
- b) a standby condition timer (see SPC-4) is enabled and that timer has expired.

The transition shall include a:

- a) Transitioning To Standby_z argument, if:
 - A) the highest priority timer that expired is the standby_z condition timer; or
 - B) the START STOP UNIT command being processed has the POWER CONDITION MODIFIER field set to 0h (i.e., standby_z power condition);

or

- b) Transitioning To Standby_y argument, if:
 - A) the highest priority timer that expired is the standby_y condition timer; or
 - B) the START STOP UNIT command being processed has the POWER CONDITION MODIFIER field set to 1h (i.e., standby_y power condition).

4.21.5.3.4 Transition SSU_PC1:Active to SSU_PC10:Wait_Stopped

This transition shall occur if:

- a) the device server processes a START STOP UNIT command (see 5.25) with the START bit set to zero and the POWER CONDITION field set to 0h (i.e., START_VALID).

4.21.5.4 SSU_PC2:Idle state

4.21.5.4.1 SSU_PC2:Idle state description

See the PC2:Idle state in SPC-4 for details about this state.

4.21.5.4.2 Transition SSU_PC2:Idle to SSU_PC4:Active_Wait

This transition shall occur if:

- a) the device server processes a START STOP UNIT command (see 5.25) with the START bit set to one and the POWER CONDITION field set to 0h (i.e., START_VALID);
- b) the device server processes a START STOP UNIT command with the POWER CONDITION field set to 1h (i.e., ACTIVE); or
- c) the device server processes a command that requires the logical unit to be in the SSU_PC1:Active state to continue processing that command.

The transition shall include a:

- a) Transitioning From Idle argument; and
- b) Transitioning From Idle_c argument if the current power condition is the idle_c power condition.

4.21.5.4.3 Transition SSU_PC2:Idle to SSU_PC5:Wait_Idle

This transition shall occur if:

- a) the following occur:
 - A) an idle condition timer is enabled and that idle condition timer has expired; and
 - B) the priority of that idle condition timer is greater than the priority of the idle condition timer associated with the current idle power condition (see SPC-4);
- or
- b) the device server processes a START STOP UNIT command (see 5.25) with the POWER CONDITION field set to 2h (i.e., IDLE) and the POWER CONDITION MODIFIER field set to a value that specifies that the logical unit transition to a lower idle power condition.

The transition shall include a:

- a) Transitioning To Idle_b argument, if:
 - A) the highest priority timer that expired is the idle_b condition timer; or
 - B) the START STOP UNIT command being processed has the POWER CONDITION MODIFIER field set to 1h (i.e., idle_b power condition);
- or
- b) Transitioning To Idle_c argument, if:
 - A) the highest priority timer that expired is the idle_c condition timer; or
 - B) the START STOP UNIT command being processed has the POWER CONDITION MODIFIER field set to 2h (i.e., idle_c power condition).

4.21.5.4.4 Transition SSU_PC2:Idle to SSU_PC6:Wait_Standby

This transition shall occur if:

- a) the device server processes a START STOP UNIT command (see 5.25) with the POWER CONDITION field set to 3h (i.e., STANDBY); or
- b) a standby condition timer is enabled and that timer has expired.

The transition shall include a:

- a) Transitioning To Standby_z argument, if:
 - A) the highest priority timer that expired is the standby_z condition timer; or
 - B) the START STOP UNIT command being processed has the POWER CONDITION MODIFIER field set to 0h (i.e., standby_z power condition);
- or
- b) Transitioning To Standby_y argument, if:
 - A) the highest priority timer that expired is the standby_y condition timer; or
 - B) the START STOP UNIT command being processed has the POWER CONDITION MODIFIER field set to 1h (i.e., standby_y power condition).

4.21.5.4.5 Transition SSU_PC2:Idle to SSU_PC7:Idle_Wait

This transition shall occur if:

- a) the device server processes a START STOP UNIT command (see 5.25) with the POWER CONDITION field set to 2h (i.e., IDLE) and the POWER CONDITION MODIFIER field set to a value that specifies that the logical unit transition to a higher idle power condition.

The transition shall include Transitioning From Idle argument and a:

- a) Transitioning To Idle_a argument, if the START STOP UNIT command being processed has the POWER CONDITION MODIFIER field set to 0h (i.e., idle_a power condition); or
- b) Transitioning To Idle_b argument, if the START STOP UNIT command being processed has the POWER CONDITION MODIFIER field set to 1h (i.e., idle_b power condition).

4.21.5.4.6 Transition SSU_PC2:Idle to SSU_PC10:Wait_Stopped

This transition shall occur if:

- a) the device server processes a START STOP UNIT command (see 5.25) with the START bit set to zero and the POWER CONDITION field set to 0h (i.e., START_VALID).

4.21.5.5 SSU_PC3:Standby state**4.21.5.5.1 SSU_PC3:Standby state description**

See the PC3:Standby state in SPC-4 for details about this state.

4.21.5.5.2 Transition SSU_PC3:Standby to SSU_PC4:Active_Wait

This transition shall occur if:

- a) the device server processes a START STOP UNIT command (see 5.25) with the START bit set to one and the POWER CONDITION field set to 0h (i.e., START_VALID);
- b) the device server processes a START STOP UNIT command with the POWER CONDITION field set to 1h (i.e., ACTIVE); or
- c) the device server processes a command that requires the logical unit to be in the SSU_PC1:Active state to continue processing that command.

The transition shall include a Transitioning From Standby argument.

4.21.5.5.3 Transition SSU_PC3:Standby to SSU_PC6:Wait_Standby

This transition shall occur if:

- a) the following occur:
 - A) the standby_z condition timer is enabled and that timer expires; and
 - B) the priority of that standby condition timer is greater than the priority of the standby condition timer associated with the current standby power condition (see SPC-4);
- or
- b) the device server processes a START STOP UNIT command with the POWER CONDITION field set to 3h (i.e., STANDBY) and the POWER CONDITION MODIFIER field set to a value that specifies that the logical unit transition to a lower standby power condition.

The transition shall include Transitioning To Standby_z argument.

4.21.5.5.4 Transition SSU_PC3:Standby to SSU_PC7:Idle_Wait

This transition shall occur if:

- a) the device server processes a START STOP UNIT command (see 5.25) with the POWER CONDITION field set to 2h (i.e., IDLE); or
- b) the device server processes a command and determines that the device server is capable of continuing the processing of that command, when the logical unit is in the SSU_PC2:Idle state.

The transition shall include a Transitioning From Standby argument and a:

- a) Transitioning To Idle_a argument, if:
 - A) the device server processes a command and determines that the device server is capable of continuing the processing of that command, when the logical unit is in the idle_a power condition;
 - or
 - B) the device server processes a START STOP UNIT command with the POWER CONDITION MODIFIER field set to 0h (i.e., idle_a power condition);
- b) Transitioning To Idle_b argument, if:
 - A) the device server processes a command and determines that the device server is capable of continuing the processing of that command, when the logical unit is in the idle_b power condition;
 - or
 - B) the device server processes a START STOP UNIT command with the POWER CONDITION MODIFIER field set to 1h (i.e., idle_b power condition);
- or
- c) Transitioning To Idle_c argument, if:
 - A) the device server processes a command and determines that the device server is capable of continuing the processing of that command, when the logical unit is in the idle_c power condition;
 - or
 - B) the device server processes a START STOP UNIT command with the POWER CONDITION MODIFIER field set to 2h (i.e., idle_c power condition).

4.21.5.5.5 Transition SSU_PC3:Standby to SSU_PC9:Standby_Wait

This transition shall occur if:

- a) the device server processes a START STOP UNIT command (see 5.25) with the POWER CONDITION field set to 3h (i.e., STANDBY) and the POWER CONDITION MODIFIER field set to a value that specifies that the logical unit transition to a higher standby power condition.

The transition shall include a Transitioning To Standby_y argument.

4.21.5.5.6 Transition SSU_PC3:Standby to SSU_PC10:Wait_Stopped

This transition shall occur if:

- a) the device server processes a START STOP UNIT command (see 5.25) with the START bit set to zero and the POWER CONDITION field set to 0h (i.e., START_VALID).

4.21.5.6 SSU_PC4:Active_Wait state

4.21.5.6.1 SSU_PC4:Active_Wait state description

While in this state:

- a) each idle condition timer that is enabled and not expired is running;
- b) each standby condition timer that is enabled and not expired is running;
- c) the device server shall provide power management pollable sense data (see SPC-4) with the sense key set to NO SENSE and the additional sense code set to LOGICAL UNIT TRANSITIONING TO ANOTHER POWER CONDITION; and
- d) the logical unit is performing the operations required for it to be in the SSU_PC1:Active state (e.g., a disk drive spins up its medium).

If this state was entered with a Transitioning From Idle argument, then:

- a) the device server is capable of processing and completing the same commands, except START STOP UNIT commands with the IMMEDIATE bit set to zero (see 5.25), that the device server is able to process and complete while in the SSU_PC2:Idle state;
- b) the peak power consumed in this state shall be no more than the typical peak power consumed in the SSU_PC1:Active state; and
- c) if:
 - A) this state was entered with a Transitioning From Idle_c argument; and
 - B) the CCF IDLE field in the Power Condition mode page (see SPC-4) is set to 10b (i.e., enabled),
 then the device server shall terminate any command, except a START STOP UNIT command, that requires the logical unit be in the SSU_PC1:Active state to continue processing, with CHECK CONDITION status, with the sense key set to NOT READY and the additional sense code set to LOGICAL UNIT IS IN PROCESS OF BECOMING READY.

If this state was entered with a Transitioning From Standby argument, then:

- a) the device server is capable of processing and completing the same commands, except START STOP UNIT commands with the IMMEDIATE bit set to zero, that the device server is able to process and complete while in the SSU_PC3:Standby state;
- b) the peak power consumption in this state is not limited by this standard; and
- c) if the CCF STANDBY field in the Power Condition mode page (see SPC-4) is set to 10b (i.e., enabled), then the device server shall terminate any command, except a START STOP UNIT command, that requires the logical unit be in the SSU_PC1:Active state or SSU_PC2:Idle state to continue processing, with CHECK CONDITION status, with the sense key set to NOT READY and the additional sense code set to LOGICAL UNIT IS IN PROCESS OF BECOMING READY.

If this state was entered with a Transitioning From Stopped argument, then:

- a) the device server is capable of processing and completing the same commands, except START STOP UNIT commands with the IMMEDIATE bit set to zero, that the device server is able to process and complete while in the SSU_PC8:Stopped state;
- b) the peak power consumption in this state is not limited by this standard; and
- c) if the CCF STOPPED field in the Power Condition mode page (see SPC-4) is set to 10b (i.e., enabled), then the device server shall terminate any TEST UNIT READY command or medium access command, with CHECK CONDITION status, with the sense key set to NOT READY and the additional sense code set to LOGICAL UNIT IS IN PROCESS OF BECOMING READY.

If this state was entered with a Transitioning From Powered On argument, then:

- a) the device server is capable of processing and completing the same commands, except START STOP UNIT commands with the IMMED bit set to zero or TEST UNIT READY command, that the device server is able to process and complete while in the SSU_PC8:Stopped state;
- b) the peak power consumption in this state is not limited by this standard; and
- c) the device server shall terminate any TEST UNIT READY command or medium access command with CHECK CONDITION status with the sense key set to NOT READY and the additional sense code set to LOGICAL UNIT IS IN PROCESS OF BECOMING READY.

If an idle condition timer or a standby condition timer is enabled and expires, then that timer is ignored in this state.

4.21.5.6.2 Transition SSU_PC4:Active_Wait to SSU_PC1:Active

See the PC4:Active_Wait to PC1:Active transition in SPC-4 for details about this transition.

4.21.5.7 SSU_PC5:Wait_Idle state

4.21.5.7.1 SSU_PC5:Wait_Idle state description

See the PC5:Wait_Idle state in SPC-4 for details about this state.

4.21.5.7.2 Transition SSU_PC5:Wait_Idle to SSU_PC2:Idle

See the PC5:Wait_Idle to PC2:Idle transition in SPC-4 for details about this transition.

4.21.5.8 SSU_PC6:Wait_Standby state

4.21.5.8.1 SSU_PC6:Wait_Standby state description

See the PC6:Wait_Standby state in SPC-4 for details about this state.

4.21.5.8.2 Transition SSU_PC6:Wait_Standby to SSU_PC3:Standby

See the PC6:Wait_Standby to PC3:Standby transition in SPC-4 for details about this transition.

4.21.5.9 SSU_PC7:Idle_Wait state

4.21.5.9.1 SSU_PC7:Idle_Wait state description

While in this state:

- a) each idle condition timer that is enabled and not expired is running;
- b) each standby condition timer that is enabled and not expired is running;
- c) the device server shall provide power management pollable sense data (see SPC-4) with the sense key set to NO SENSE and the additional sense code set to LOGICAL UNIT TRANSITIONING TO ANOTHER POWER CONDITION; and
- d) the logical unit is performing the operations required for it to be in the SSU_PC2:Idle state (e.g., a disk drive spins up its medium).

If this state was entered with a Transitioning From Idle argument, then:

- a) the device server is capable of processing and completing the same commands, except START STOP UNIT commands with the IMMED bit set to zero (see 5.25), that the device server is able to process and complete while in the SSU_PC2:Idle state; and
- b) the peak power consumed in this state shall be no more than the typical peak power consumed in the SSU_PC1:Active state.

If this state was entered with a Transitioning From Standby argument, then:

- a) the device server is capable of processing and completing the same commands, except START STOP UNIT commands with the IMMED bit set to zero, that the device server is able to process and complete while in the SSU_PC3:Standby state;
- b) the peak power consumption in this state is not limited by this standard; and
- c) the CCF STANDBY field in the Power Condition mode page (see SPC-4) is set to 10b (i.e., enabled), then the device server shall terminate any command, except a START STOP UNIT command, that requires the logical unit be in the SSU_PC1:Active state or SSU_PC2:Idle state to continue processing, with CHECK CONDITION status, with the sense key set to NOT READY and the additional sense code set to LOGICAL UNIT IS IN PROCESS OF BECOMING READY.

If this state was entered with a Transitioning From Stopped argument, then:

- a) the device server is capable of processing and completing the same commands, except START STOP UNIT commands with the IMMED bit set to zero, that the device server is able to process and complete while in the SSU_PC8:Stopped state;
- b) the peak power consumption in this state is not limited by this standard; and
- c) if the CCF STOPPED field in the Power Condition mode page (see SPC-4) is set to 10b (i.e., enabled), then the device server shall terminate any TEST UNIT READY command or medium access command, with CHECK CONDITION status, with the sense key set to NOT READY and the additional sense code set to LOGICAL UNIT IS IN PROCESS OF BECOMING READY.

If an idle condition timer or a standby condition timer is enabled and expires, then that timer is ignored in this state.

4.21.5.9.2 Transition SSU_PC7:Idle_Wait to SSU_PC2:Idle

This transition shall occur when the logical unit meets the requirements for being in the:

- a) idle_a power condition, if this state was entered with a Transitioning To Idle_a argument;
- b) idle_b power condition, if this state was entered with a Transitioning To Idle_b argument; or
- c) idle_c power condition, if this state was entered with a Transitioning To Idle_c argument.

4.21.5.10 SSU_PC8:Stopped state

4.21.5.10.1 SSU_PC8:Stopped state description

While in this state:

- a) the logical unit is in the stopped power condition (see 4.21.4);
- b) the idle condition timers and the standby condition timers are disabled;
- c) the device server shall provide power management pollable sense data (see SPC-4); and
- d) the device server shall terminate each medium access command or TEST UNIT READY command (see SPC-4) as described in 4.21.4.

4.21.5.10.2 Transition SSU_PC8:Stopped to SSU_PC4:Active_Wait

This transition shall occur if:

- a) the device server processes a START STOP UNIT command (see 5.25) with the START bit set to one and the POWER CONDITION field set to 0h (i.e., START_VALID); or
- b) the device server processes a START STOP UNIT command with the POWER CONDITION field set to 1h (i.e., ACTIVE).

The transition shall include a Transitioning From Stopped argument.

4.21.5.10.3 Transition SSU_PC8:Stopped to SSU_PC7:Idle_Wait

This transition shall occur if:

- a) the device server processes a START STOP UNIT command (see 5.25) with the POWER CONDITION field set to 2h (i.e., IDLE).

The transition shall include a Transitioning From Stopped argument and a:

- a) Transitioning To Idle_a argument, if the START STOP UNIT command being processed has the POWER CONDITION MODIFIER field set to 0h (i.e., idle_a power condition);
- b) Transitioning To Idle_b argument, if the START STOP UNIT command being processed has the POWER CONDITION MODIFIER field set to 1h (i.e., idle_b power condition); or
- c) Transitioning To Idle_c argument, if the START STOP UNIT command being processed has the POWER CONDITION MODIFIER field set to 2h (i.e., idle_c power condition).

4.21.5.10.4 Transition SSU_PC8:Stopped to SSU_PC9:Standby_Wait

This transition shall occur if:

- a) the device server processes a START STOP UNIT command (see 5.25) with the POWER CONDITION field set to STANDBY.

The transition shall include a Transitioning From Stopped argument and a:

- a) Transitioning To Standby_z argument, if the START STOP UNIT command being processed has the POWER CONDITION MODIFIER field set to 0h (i.e., standby_z power condition); or
- b) Transitioning To Standby_y argument, if the START STOP UNIT command being processed has the POWER CONDITION MODIFIER field set to 1h (i.e., standby_y power condition).

4.21.5.11 SSU_PC9:Standby_Wait state

4.21.5.11.1 SSU_PC9:Standby_Wait state description

While in this state:

- a) the device server shall provide power management pollable sense data (see SPC-4) with the sense key set to NO SENSE and the additional sense code set to LOGICAL UNIT TRANSITIONING TO ANOTHER POWER CONDITION;
- b) the peak power consumed in this state shall be no more than the typical peak power consumed in the SSU_PC1:Active state; and
- c) the logical unit is performing the operations required for it to be in the SSU_PC3:Standby state ((e.g., a direct access block device is activating circuitry).

If this state was entered with a Transitioning From Standby argument, then the device server is capable of processing and completing the same commands, except START STOP UNIT commands with the IMMED bit set to zero (see 5.25), that the device server is able to process and complete in the SSU_PC3:Standby state.

If this state was entered with a Transitioning From Stopped argument, then:

- a) the device server is capable of processing and completing the same commands, except START STOP UNIT commands with the IMMED bit set to zero, that the device server is able to process and complete while in the SSU_PC8:Stopped state; and
- b) if the CCF STOPPED field in the Power Condition mode page (see SPC-4) is set to 10b (i.e., enabled), then the device server shall terminate any TEST UNIT READY command or medium access command, with CHECK CONDITION status, with the sense key set to NOT READY and the additional sense code set to LOGICAL UNIT IS IN PROCESS OF BECOMING READY.

4.21.5.11.2 Transition SSU_PC9:Standby_Wait to SSU_PC3:Standby

This transition shall occur when the logical unit meets the requirements for being in the:

- a) standby_y power condition, if this state was entered with a Transitioning To Standby_y argument; or

- b) standby_z power condition, if this state was entered with a Transitioning To Standby_z argument.

4.21.5.12 SSU_PC10:Wait_Stopped state

4.21.5.12.1 SSU_PC10:Wait_Stopped state description

While in this state:

- a) the device server shall provide power management pollable sense data (see SPC-4) with the sense key set to NO SENSE and the additional sense code set to LOGICAL UNIT TRANSITIONING TO ANOTHER POWER CONDITION;
- b) the device server is capable of processing and completing the same commands, except START STOP UNIT commands with the IMMED bit set to zero (see 5.25), that the device server is able to process and complete in the SSU_PC8:Stopped state;
- c) the logical unit is performing the operations required for it to be in the SSU_PC8:Stopped state (e.g., a disk drive spins down its medium); and
- d) the device server shall terminate any TEST UNIT READY command or medium access command with CHECK CONDITION status with the sense key set to NOT READY and the additional sense code set to LOGICAL UNIT NOT READY, INITIALIZING COMMAND REQUIRED.

4.21.5.12.2 Transition SSU_PC10:Wait_Stopped to SSU_PC8:Stopped

This transition shall occur when:

- a) the logical unit meets the requirements for being in the SSU_PC8:Stopped state.

4.22 Protection information model

4.22.1 Protection information overview

The protection information model provides for protection of user data while user data is being transferred between a sender and a receiver. Protection information is generated at the application layer and may be checked by any object associated with the I_T_L nexus (see SAM-5). Once received, protection information is retained (e.g., written to the medium, stored in non-volatile memory, or recalculated on read back) by the device server until overwritten. Power loss, hard reset, logical unit reset, and I_T nexus loss shall have no effect on the retention of protection information.

Support for protection information shall be indicated in the PROTECT bit in the standard INQUIRY data (see SPC-4).

If the logical unit is formatted with protection information, and the EMDP bit is set to one in the Disconnect-Reconnect mode page (see SPC-4), then checking of the logical block reference tag within a service delivery subsystem without accounting for modified data pointers and data alignments may cause false errors when logical blocks are transmitted out of order.

Protection information is also referred to as the data integrity field (DIF).

4.22.2 Protection types

4.22.2.1 Protection types overview

The content of protection information is dependent on the type of protection to which a logical unit has been formatted.

The type of protection supported by the logical unit shall be indicated in the SPT field in the Extended INQUIRY Data VPD page (see SPC-4). The current protection type shall be indicated in the P_TYPE field in the READ CAPACITY (16) parameter data (see 5.16.2).

An application client may format the logical unit to a specific type of protection using the FMTPINFO field and the PROTECTION FIELD USAGE field in the FORMAT UNIT command (see 5.3).

An application client may format the logical unit to place protection information at intervals other than on logical block boundaries using the PROTECTION INTERVAL EXPONENT field in the FORMAT UNIT command.

A medium access command is processed in a different manner by a device server depending on the type of protection in effect. When used in relation to types of protection, the term “medium access command” is defined as any one of the following commands:

- a) COMPARE AND WRITE;
- b) ORWRITE (16);
- c) ORWRITE (32);
- d) READ (10);
- e) READ (12);
- f) READ (16);
- g) READ (32);
- h) VERIFY (10);
- i) VERIFY (12);
- j) VERIFY (16);
- k) VERIFY (32);
- l) WRITE (10);
- m) WRITE (12);
- n) WRITE (16);
- o) WRITE (32);
- p) WRITE AND VERIFY (10);
- q) WRITE AND VERIFY (12);
- r) WRITE AND VERIFY (16);
- s) WRITE AND VERIFY (32);
- t) WRITE SAME (10);
- u) WRITE SAME (16);
- v) WRITE SAME (32);
- w) XDWRITEREAD (10); and
- x) XDWRITEREAD (32).

4.22.2.2 Type 0 protection

Type 0 protection defines no protection over that which is defined within the transport protocol.

A logical unit that has been formatted with protection information disabled (see 5.2) or a logical unit that does not support protection information (i.e., the PROTECT bit set to zero in the standard INQUIRY data (see SPC-4)) has type 0 protection.

If type 0 protection is enabled and the RDPROTECT field, the WRPROTECT field, the VRPROTECT field, or the ORPROTECT field is set to a non-zero value, then medium access commands are invalid and may be terminated by the device server with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

If type 0 protection is enabled, then the following medium access commands are invalid and shall be terminated by the device server with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID COMMAND OPERATION CODE:

- a) READ (32);
- b) VERIFY (32);
- c) WRITE (32);
- d) WRITE AND VERIFY (32); and
- e) WRITE SAME (32).

4.22.2.3 Type 1 protection

Type 1 protection:

- a) defines the content of each LOGICAL BLOCK GUARD field;
- b) does not define the content of any LOGICAL BLOCK APPLICATION TAG field; and
- c) defines the content of each LOGICAL BLOCK REFERENCE TAG field.

If type 1 protection is enabled, then the following medium access commands are invalid and shall be terminated by the device server with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID COMMAND OPERATION CODE:

- a) READ (32);
- b) VERIFY (32);
- c) WRITE (32);
- d) WRITE AND VERIFY (32); and
- e) WRITE SAME (32).

For valid medium access commands in which the RDPROTECT field, the WRPROTECT field, the VRPROTECT field, or the ORPROTECT field is set to:

- a) zero, the Data-In Buffer and/or Data-Out Buffer associated with those commands shall consist of logical block data containing only user data; or
- b) a non-zero value, the Data-In Buffer and/or Data-Out Buffer shall consist of logical block data containing both user data and protection information.

4.22.2.4 Type 2 protection

Type 2 protection:

- a) defines the content of each LOGICAL BLOCK GUARD field;
- b) does not define the content of any LOGICAL BLOCK APPLICATION TAG field; and
- c) defines, except for the first logical block addressed by the command, the content of each LOGICAL BLOCK REFERENCE TAG field.

If type 2 protection is enabled and the RDPROTECT field, the WRPROTECT field, the VRPROTECT field, or the ORPROTECT field is set to a non-zero value, then the following medium access commands are invalid and shall be terminated by the device server with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID COMMAND OPERATION CODE:

- a) COMPARE AND WRITE;
- b) ORWRITE (16);
- c) ORWRITE (32);
- d) READ (10);
- e) READ (12);
- f) READ (16);
- g) VERIFY (10);
- h) VERIFY (12);
- i) VERIFY (16);
- j) WRITE (10);
- k) WRITE (12);
- l) WRITE (16);
- m) WRITE AND VERIFY (10);
- n) WRITE AND VERIFY (12);
- o) WRITE AND VERIFY (16);
- p) WRITE SAME (10);
- q) WRITE SAME (16);
- r) XDWRITEREAD (10); and
- s) XDWRITEREAD (32).

For valid medium access commands in which the RDPROTECT field, the WRPROTECT field, the VRPROTECT field, or the ORPROTECT field is set to:

- a) zero, the Data-In Buffer and/or Data-Out Buffer associated with those commands shall consist of logical block data containing only user data; or
- b) a non-zero value, the Data-In Buffer and/or Data-Out Buffer shall consist of logical block data containing both user data and protection information.

4.22.2.5 Type 3 protection

Type 3 protection:

- a) defines the content of each LOGICAL BLOCK GUARD field;
- b) does not define the content of any LOGICAL BLOCK APPLICATION TAG field; and
- c) does not define the content of any LOGICAL BLOCK REFERENCE TAG field.

If type 3 protection is enabled, then the following medium access commands are invalid and shall be terminated by the device server with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID COMMAND OPERATION CODE:

- a) READ (32);
- b) VERIFY (32);
- c) WRITE (32);
- d) WRITE AND VERIFY (32); and
- e) WRITE SAME (32).

For valid medium access commands in which the RDPROTECT field, the WRPROTECT field, the VRPROTECT field, or the ORPROTECT field is set to:

- a) zero, the Data-In Buffer and/or Data-Out Buffer associated with those commands shall consist of logical block data containing only user data; or
- b) a non-zero value, the Data-In Buffer and/or Data-Out Buffer shall consist of logical block data containing both user data and protection information.

4.22.3 Protection information format

Table 20 defines the placement of protection information in a logical block with a single protection information interval (i.e., the PROTECTION INTERVAL EXPONENT field is set to zero in the parameter list header for a FORMAT UNIT command (see 5.3.2.2))

Table 20 – Logical block data format with a single protection information interval

Byte	Bit	7	6	5	4	3	2	1	0
0		USER DATA							
...									
n - 1									
n	(MSB)	LOGICAL BLOCK GUARD							
n + 1		(LSB)							
n + 2	(MSB)	LOGICAL BLOCK APPLICATION TAG							
n + 3									
n + 4	(MSB)	LOGICAL BLOCK REFERENCE TAG							
...									
n + 7									

Table 21 shows an example of the placement of protection information in a logical block with more than one protection information interval (i.e., the PROTECTION INTERVAL EXPONENT field is set to a non-zero value in the parameter list header for a FORMAT UNIT command (see 5.3.2.2)).

Table 21 – An example of the logical block data for a logical block with more than one protection information interval

Byte	Bit	7	6	5	4	3	2	1	0
0		USER DATA [first]							
...									
n - 1									
n	(MSB)	LOGICAL BLOCK GUARD [first]							
n + 1									
n + 2	(MSB)	LOGICAL BLOCK APPLICATION TAG [first]							
n + 3									
n + 4	(MSB)	LOGICAL BLOCK REFERENCE TAG [first]							
...									
n + 7									
n + 8		USER DATA [second]							
...									
m - 1									
m	(MSB)	LOGICAL BLOCK GUARD [second]							
m + 1									
m + 2	(MSB)	LOGICAL BLOCK APPLICATION TAG [second]							
m + 3									
m + 4	(MSB)	LOGICAL BLOCK REFERENCE TAG [second]							
...									
m + 7									
		⋮							
...		USER DATA [last]							
z - 1									
z	(MSB)								
z + 1		(LSB)							
z + 2	(MSB)	LOGICAL BLOCK APPLICATION TAG [last]							
z + 3									
z + 4	(MSB)	LOGICAL BLOCK REFERENCE TAG [last]							
...									
z + 7									

Each USER DATA field shall contain user data.

Each LOGICAL BLOCK GUARD field contains a CRC (see 4.22.4). Only the contents of the USER DATA field immediately preceding THE LOGICAL BLOCK GUARD field (i.e., the user data between the preceding logical block

reference tag, if any, and the current logical block guard) shall be used to generate and check the CRC contained in the LOGICAL BLOCK GUARD field.

Each LOGICAL BLOCK APPLICATION TAG field is set by the application client. If the device server detects a:

- a) LOGICAL BLOCK APPLICATION TAG field set to FFFFh and type 1 protection (see 4.22.2.3) is enabled;
- b) LOGICAL BLOCK APPLICATION TAG field set to FFFFh and type 2 protection (see 4.22.2.4) is enabled; or
- c) LOGICAL BLOCK APPLICATION TAG field set to FFFFh, LOGICAL BLOCK REFERENCE TAG field set to FFFF_FFFFh, and type 3 protection (see 4.22.2.5) is enabled,

then the device server disables checking of all protection information for the associated protection information interval when performing a read operation. Otherwise, if the ATMPE bit in the Control mode page (see SPC-4) is:

- a) set to one, then the logical block application tags are defined by the Application Tag mode page (see 6.5.3); or
- b) set to zero, then the logical block application tags are not defined by this standard.

The LOGICAL BLOCK APPLICATION TAG field may be modified by a device server if the ATO bit is set to zero in the Control mode page (see SPC-4). If the ATO bit is set to one in the Control mode page, then the device server shall not modify the LOGICAL BLOCK APPLICATION TAG field.

The contents of a LOGICAL BLOCK APPLICATION TAG field shall not be used to generate or check the CRC contained in the LOGICAL BLOCK GUARD field.

The first LOGICAL BLOCK REFERENCE TAG field of the first logical block in the Data-In Buffer and/or Data-Out Buffer shall contain the value specified in table 22.

Table 22 – Content of the first LOGICAL BLOCK REFERENCE TAG field for the first logical block in the Data-In Buffer and/or Data-Out Buffer

Protection Type	Content of the first LOGICAL BLOCK REFERENCE TAG field for the first logical block in the Data-In Buffer and/or Data-Out Buffer
Type 1 ^a protection (see 4.22.2.3)	The least significant four bytes of the LBA contained in the LOGICAL BLOCK ADDRESS field of the CDB.
Type 2 protection (see 4.22.2.4)	The value in the EXPECTED INITIAL LOGICAL BLOCK REFERENCE TAG field of the CDB.
Type 3 protection (see 4.22.2.5)	Not defined in this standard. If the ATO bit is set to zero in the Control mode page (see SPC-4), then this field may be modified by the device server. If the ATO bit is set to one in the Control mode page, then the device server shall not modify this field.
^a The length of the protection information interval is equal to the logical block length (see 5.3.2).	

Subsequent LOGICAL BLOCK REFERENCE TAG fields for a logical block in the Data-In Buffer and/or Data-Out Buffer shall be set as specified in table 23.

Table 23 – Content of subsequent LOGICAL BLOCK REFERENCE TAG fields for a logical block in the Data-In Buffer and/or Data-Out Buffer

Protection Type	The content of subsequent LOGICAL BLOCK REFERENCE TAG fields in the Data-In Buffer and/or Data-Out Buffer
Type 1 protection (see 4.22.2.3) and Type 2 protection (see 4.22.2.4)	The previous logical block reference tag plus one. If the contents of the previous LOGICAL BLOCK REFERENCE TAG field is FFFF_FFFFh, then the contents of the subsequent LOGICAL BLOCK REFERENCE TAG field is 0000_0000h.
Type 3 protection (see 4.22.2.5)	Not defined in this standard. If the ATO bit is set to zero in the Control mode page (see SPC-4), then this field may be modified by the device server. If the ATO bit is set to one in the Control mode page, then the device server shall not modify this field.

The contents of a LOGICAL BLOCK REFERENCE TAG field shall not be used to generate or check the CRC contained in the LOGICAL BLOCK GUARD field.

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC 14776-323:2017

4.22.4 Logical block guard

4.22.4.1 Logical block guard overview

A LOGICAL BLOCK GUARD field shall contain a CRC that is generated from the contents of only the USER DATA field immediately preceding the LOGICAL BLOCK GUARD field.

Table 24 defines the CRC polynomials used to generate the logical block guard from the contents of the USER DATA field.

Table 24 – CRC polynomials

Function	Definition
F(x)	A polynomial representing the transmitted USER DATA field, which is covered by the CRC. For the purposes of the CRC, the coefficient of the highest order term shall be byte zero bit seven of the USER DATA field and the coefficient of the lowest order term shall be bit zero of the last byte of the USER DATA field.
F'(x)	A polynomial representing the received USER DATA field.
G(x)	The generator polynomial: $G(x) = x^{16} + x^{15} + x^{11} + x^9 + x^8 + x^7 + x^5 + x^4 + x^2 + x + 1$ (i.e., in finite field notation $G(x) = 1_8BB7h$)
R(x)	The remainder polynomial calculated during CRC generation by the transmitter, representing the transmitted LOGICAL BLOCK GUARD field.
R'(x)	A polynomial representing the received LOGICAL BLOCK GUARD field.
RB(x)	The remainder polynomial calculated during CRC checking by the receiver. $RB(x) = 0$ indicates no error was detected.
RC(x)	The remainder polynomial calculated during CRC checking by the receiver. $RC(x) = 0$ indicates no error was detected.
QA(x)	The quotient polynomial calculated during CRC generation by the transmitter. The value of QA(x) is not used.
QB(x)	The quotient polynomial calculated during CRC checking by the receiver. The value of QB(x) is not used.
QC(x)	The quotient polynomial calculated during CRC checking by the receiver. The value of QC(x) is not used.
M(x)	A polynomial representing the transmitted USER DATA field followed by the transmitted LOGICAL BLOCK GUARD field.
M'(x)	A polynomial representing the received USER DATA field followed by the received LOGICAL BLOCK GUARD field.

4.22.4.2 CRC generation

The equations that are used to generate the CRC from F(x) are as follows. All arithmetic is modulo 2.

The transmitter shall calculate the CRC by appending 16 zeros to F(x) and dividing by G(x) to obtain the remainder R(x):

$$\frac{(x^{16} \times F(x))}{G(x)} = QA(x) + \frac{R(x)}{G(x)}$$

R(x) is the CRC value, and is transmitted in the LOGICAL BLOCK GUARD field.

$M(x)$ is the polynomial representing the USER DATA field followed by the LOGICAL BLOCK GUARD field (i.e., $F(x)$ followed by $R(x)$):

$$M(x) = (x^{16} \times F(x)) + R(x)$$

4.22.4.3 CRC checking

$M'(x)$ (i.e., the polynomial representing the received USER DATA field followed by the received LOGICAL BLOCK GUARD field) may differ from $M(x)$ (i.e., the polynomial representing the transmitted USER DATA field followed by the transmitted LOGICAL BLOCK GUARD field) if there are transmission errors.

The receiver may check $M'(x)$ validity by appending 16 zeros to $F'(x)$ and dividing by $G(x)$ and comparing the calculated remainder $RB(x)$ to the received CRC value $R'(x)$:

$$\frac{(x^{16} \times F'(x))}{G(x)} = QB(x) + \frac{RB(x)}{G(x)}$$

In the absence of errors in $F'(x)$ and $R'(x)$, the remainder $RB(x)$ is equal to $R'(x)$.

The receiver may check $M'(x)$ validity by dividing $M'(x)$ by $G(x)$ and comparing the calculated remainder $RC(x)$ to zero:

$$\frac{M'(x)}{G(x)} = QC(x) + \frac{RC(x)}{G(x)}$$

In the absence of errors in $F'(x)$ and $R'(x)$, the remainder $RC(x)$ is equal to zero.

Both methods of checking $M'(x)$ validity are mathematically equivalent.

4.22.4.4 CRC test cases

Several CRC test cases are shown in table 25.

Table 25 – CRC test cases

Pattern	CRC
32 bytes each set to 00h	0000h
32 bytes each set to FFh	A293h
32 bytes of an incrementing pattern from 00h to 1Fh	0224h
2 bytes each set to FFh followed by 30 bytes set to 00h	21B8h
32 bytes of a decrementing pattern from FFh to E0h	A0B7h

4.22.5 Application of protection information

Before an application client transmits or receives logical block data with protection information, the application client:

- 1) determines if a logical unit supports protection information using the INQUIRY command (see the PROTECT bit in the standard INQUIRY data in SPC-4);
- 2) if protection information is supported, then determines if the logical unit is formatted to accept protection information using the READ CAPACITY (16) command (e.g., see the PROT_EN bit and the P_TYPE field in the returned parameter data (see 5.16.2)); and
- 3) if the logical unit supports protection information and is not formatted to accept protection information, then formats the logical unit (see 5.3) with protection information enabled.

If the logical unit supports protection information and is formatted to accept protection information, then the application client may use read commands that support protection information and should use verify commands and write commands that support protection information.

4.22.6 Protection information and commands

The enabling of protection information enables fields in medium access commands that instruct the device server on the handling of protection information. The detailed definitions of each command's protection information fields are in the individual command descriptions.

The commands that are affected while protection information is enabled are listed in table 31.

Commands that cause a device server to return the length in bytes of each logical block (e.g., the MODE SENSE (10) command and the READ CAPACITY (16) command) shall cause the device server to return the combined length of the USER DATA field(s) contained in the logical block, not including the length of any protection information (i.e., the LOGICAL BLOCK GUARD field(s), the LOGICAL BLOCK APPLICATION TAG field(s), and the LOGICAL BLOCK REFERENCE TAG field(s)) (e.g., if the user data plus the protection information is equal to 520 bytes and there is one protection information interval, then 512 is returned).

4.23 Grouping function

A grouping function is a function that collects information about attributes associated with commands (i.e., information about commands with the same group value are collected into the specified group). The definition of the attributes and the groups is outside the scope of this standard. Groups are identified with the GROUP NUMBER field in the CDBs of certain commands (e.g., the PRE-FETCH (10) command (see 5.8)).

Support for the grouping function is indicated in the GROUP_SUP bit in the Extended INQUIRY Data VPD page (see SPC-4).

The collection of this information is outside the scope of this standard (e.g., the information may not be transmitted using any SCSI protocols).

EXAMPLE - In a SCSI domain in which two applications are using a subsystem where one application streams data and another accesses data randomly, if:

- a) the streaming application groups all of its commands with one group number (e.g., x); and
- b) the random application groups all of its commands with another group number (e.g., y),

then the applications use those group numbers (e.g., x and y) to collect separate performance metrics for each application.

A management application then reads the performance metrics and determines if the performance of a specific group is acceptable.

4.24 Background scan operations

4.24.1 Background scan overview

A background scan operation is either a background pre-scan operation (see 4.24.2) or a background medium scan operation (see 4.24.3).

During a background scan operation, the device server performs read medium operations for the purpose of:

- a) identifying logical blocks that are difficult to read (i.e., recoverable) or unreadable (i.e., unrecoverable);
- b) logging problems encountered during the background scan operation; and
- c) when allowed, taking a vendor specific action to repair recoverable logical blocks or perform automatic read reallocation of recoverable logical blocks.

During a background scan operation, if a read medium operation encounters a recovered error (i.e., a logical block is readable but requires extra actions (e.g., retries or application of a correction algorithm) to be read),

then the device server may resolve the problem using vendor specific means. The value of the ARRE bit in the Read-Write Error Recovery mode page (see 6.5.8) determines whether or not the device server performs automatic read reassignment.

During a background scan operation, if a read medium operation encounters an unrecovered error (i.e., a logical block is unreadable), then the device server may mark the logical block unrecoverable. The value of the AWRE bit in the Read-Write Error Recovery mode page (see 6.5.8) determines whether or not the device server performs automatic write reassignment. If the AWRE bit is set to one, then the device server performs automatic write reassignment at the start of the next write medium operation accessing that logical block.

During a background scan operation, the device server:

- a) may scan the logical blocks in any order (e.g., based on physical block layout);
- b) should not retain any data from logical blocks in cache memory after the logical blocks are read;
- c) shall ignore pseudo unrecovered errors with correction disabled (see 4.18.2); and
- d) shall process pseudo unrecovered errors with correction enabled.

4.24.2 Background pre-scan operations

4.24.2.1 Enabling background pre-scan operations

A background pre-scan operation is enabled after:

- 1) the EN_PS bit in the Background Control mode page (see 6.5.4) is set to zero;
 - 2) the EN_PS bit is set to one; and
 - 3) the SCSI device is power cycled if;
 - A) the S_L_FULL bit in the Background Control mode page is:
 - a) set to zero; or
 - b) set to one and the Background Scan log parameters in the Background Scan Results log page (see 6.4.2) are not all used;
- and
- B) the saved value of the EN_PS bit is set to one.

After a background pre-scan operation is enabled, the device server shall:

- a) initialize the Background Pre-scan Time Limit timer to the time specified in the BACKGROUND PRE-SCAN TIME LIMIT field in the Background Control mode page and start the timer;
- b) initialize the Background Medium Scan Interval timer to the time specified in the BACKGROUND MEDIUM SCAN INTERVAL TIME field in the Background Control mode page and start the timer; and
- c) begin the background pre-scan operation (i.e., begin scanning the medium).

4.24.2.2 Suspending and resuming background pre-scan operations

A background pre-scan operation shall be suspended when any of the following occurs:

- a) a command or task management function is processed that requires the background pre-scan operation to be suspended;
- b) a SCSI event (e.g., a hard reset) (see SAM-5) is processed that requires the background pre-scan operation to be suspended;
- c) a power condition timer expires (see the Power Condition mode page in SPC-4), and the PM_BG_PRECEDENCE field in the Power Condition mode page is set to 10b; or
- d) the S_L_FULL bit in the Background Control mode page (see 6.5.4) is set to one, and the Background Scan log parameters in the Background Scan Results log page (see 6.4.2) are all used.

If a command is received that requires a background pre-scan operation to be suspended, then the following should occur within the time specified in the MAXIMUM TIME TO SUSPEND BACKGROUND SCAN field in the Background Control mode page:

- a) the logical unit suspends the background medium scan operation; and
- b) the device server begins processing the command.

If a background pre-scan operation is suspended, then the device server shall not stop:

- a) the Background Pre-scan Time Limit timer;
- b) the Background Medium Scan Interval timer; and
- c) any process that results in an event that causes a background function to occur (e.g., not stop any timers or counters associated with background functions).

While a background pre-scan operation is suspended and not halted (see 4.24.3.2), the device server shall convert each write operation accessing a logical block that has not been scanned during the background pre-scan operation into a write medium operation followed by a verify medium operation in order to verify that the logical block data just written was read back without error. If a write medium operation accesses a logical block that has already been scanned during the background pre-scan operation, then the device server shall not perform the additional verify medium operation.

A background pre-scan operation shall be resumed from where the operation was suspended when:

- a) there are no commands in any task set to be processed;
- b) there are no task management functions to be processed;
- c) there are no SCSI events to be processed;
- d) no ACA condition exists;
- e) the PM_BG_PRECEDENCE field in the Power Condition mode page is set to 10b (see SPC-4), but no power condition timer defined in the Power Condition mode page has expired;
- f) the S_L_FULL bit in the Background Control mode page is set to zero, or the Background Medium Scan log parameters in the Background Scan Results log page are not all used;
- g) the logical unit has been idle for the time specified in the MINIMUM IDLE TIME BEFORE BACKGROUND SCAN field in the Background Control mode page; and
- h) the background pre-scan operation has not been halted (see 4.24.3.2).

4.24.2.3 Halting background pre-scan operations

The device server shall halt a background pre-scan operation if any of the following occurs:

- a) the background pre-scan operation completes scanning all logical blocks on the medium;
- b) an application client sets the EN_PS bit to zero in the Background Control mode page (see 6.5.4);
- c) the Background Pre-scan Time Limit timer expires;
- d) the device server detects a fatal error;
- e) the device server detects a vendor specific pattern of errors;
- f) the device server detects a medium formatted without a PLIST (see 4.13); or
- g) the device server detects temperature out of range.

After a background pre-scan operation has been halted, the device server shall not enable a background operation until the conditions in 4.24.2.1 are met.

4.24.3 Background medium scan

4.24.3.1 Enabling background medium scan operations

Background medium scan operations are enabled if:

- a) a background pre-scan operation (see 4.24.2) is not in progress;
- b) the S_L_FULL bit in the Background Control mode page (see 6.5.4) is:
 - A) set to zero; or
 - B) set to one and the Background Scan log parameters in the Background Scan Results log page (see 6.4.2) are not all used;

and

- c) the EN_BMS bit in the Background Control mode page is set to one.

If background medium scan operations are enabled, then the device server shall begin a background medium scan operation (i.e., begin scanning the medium) when:

- a) the Background Medium Scan Interval timer has expired; and

- b) the logical unit has been idle for the time specified in the MINIMUM IDLE TIME BEFORE BACKGROUND SCAN field in the Background Control mode page.

After power on, if background pre-scan operations are not enabled (see 4.24.2.1), then the device server shall set the Background Medium Scan Interval timer to zero (i.e., expired).

Whenever a background medium scan operation begins, the device server shall set the Background Medium Scan Interval timer to the time specified in the BACKGROUND MEDIUM SCAN INTERVAL TIME field in the Background Control mode page and start the timer.

4.24.3.2 Suspending and resuming background medium scan operations

The logical unit shall suspend a background medium scan operation if any of the following occurs:

- a) a command or task management function is processed that requires the background medium scan operation to be suspended;
- b) a SCSI event (e.g., a hard reset) (see SAM-5) is processed that requires the background medium scan operation to be suspended;
- c) a power condition timer expires (see the Power Condition mode page in SPC-4), and the PM_BG_PRECEDENCE field in the Power Condition mode page is set to 10b;
- d) the S_L_FULL bit in the Background Control mode page (see 6.5.4) is set to one, and the Background Scan log parameters in the Background Scan Results log page (see 6.4.2) are all used; or
- e) an application client sets the EN_BMS bit in the Background Control mode page to zero.

If a command is received that requires a background medium scan operation to be suspended, then the following should occur within the time specified in the MAXIMUM TIME TO SUSPEND BACKGROUND SCAN field in the Background Control mode page:

- a) the logical unit suspends the background medium scan operation; and
- b) and the device server begins processing the command.

If a background pre-scan operation is suspended, then the device server shall not stop:

- a) the Background Medium Scan Interval timer; and
- b) any process that results in an event that causes a background function to occur (e.g., not stop any timers or counters associated with background functions).

The logical unit shall resume a suspended background medium scan operation from where the operation was suspended when:

- a) there are no commands in any task set to be processed;
- b) there are no task management functions to be processed;
- c) there are no SCSI events to be processed;
- d) the PM_BG_PRECEDENCE field in the Power Condition mode page is set to 10b (see SPC-4), but no power condition timer defined in the Power Condition mode page has expired;
- e) the S_L_FULL bit in the Background Control mode page is set to zero, or the Background Medium Scan log parameters in the Background Scan Results log page are not all used;
- f) the EN_BMS bit in the Background Control mode page is set to one; and
- g) the logical unit has been idle for the time specified in the MINIMUM IDLE TIME BEFORE BACKGROUND SCAN field in the Background Control mode page.

4.24.3.3 Halting background medium scan operations

The device server shall halt background medium scan operations if any of the following occurs:

- a) the background medium scan operation completes scanning all logical blocks on the medium;
- b) the device server detects a fatal error;
- c) the device server detects a vendor specific pattern of errors;
- d) the device server detects a medium formatted without a PLIST (see 4.13); or
- e) the device server detects temperature out of range.

After background medium scan operations have been halted, the device server shall not enable a background medium scan operation until the conditions in 4.24.3.1 are met.

4.24.4 Interpreting the logged background scan results

An application client may:

- a) poll the Background Scan Results log page (see 6.4.2) to get information about background pre-scan and background medium scan activity; or
- b) use the EBACKERR bit and the MRIE field in the Informational Exceptions Control mode page (see 6.5.6) to select a method of indicating that a medium error was detected.

If the EBACKERR bit is set to one and a medium error was detected, then the device server shall return the following additional sense codes using the method defined by the value in the MRIE field:

- a) WARNING - BACKGROUND PRE-SCAN DETECTED MEDIUM ERROR, if the failure occurs during a background pre-scan operation; or
- b) WARNING - BACKGROUND MEDIUM SCAN DETECTED MEDIUM ERROR, if the failure occurs during a background medium scan operation.

The Background Scan Status log parameter (see 6.4.2.2) in the Background Scan Results log page (see 6.4.2) indicates:

- a) whether or not a background scan operation is active or halted;
- b) the number of background scan operations that have been performed on the medium; and
- c) the progress of a background scan operation, if active.

This information may be used by an application client to monitor the background scan operations and should be used by an application client after notification via an informational exception (see 6.5.6).

The Background Scan Results log parameters (see 6.4.2.3), if any, in the Background Scan Results log page describe the LBA and the reassignment status of each logical block that generated recovered errors or unrecovered errors during the background scan's read medium operations.

After an application client analyzes the Background Scan Results log parameters and has completed actions, if any, to repair any of the indicated LBAs, the application client may delete all Background Scan Results parameters by issuing a LOG SELECT command (e.g., with the PCR bit set to one in the CDB or with the PC field set to 11b and the PARAMETER LIST LENGTH field set to zero in the CDB) (see SPC-4).

A background medium scan operation may continue to run during log page accesses. To ensure that the values in the Background Scan Results log page do not change during a sequence of accesses, the application client:

- 1) sets the EN_BMS bit to zero in the Background Control mode page in order to suspend the background medium scan operation;
- 2) reads the Background Scan Results log page with a LOG SENSE command;
- 3) processes the Background Scan Results log page;
- 4) deletes the Background Scan Results log page entries with the LOG SELECT command (e.g., with the PCR bit set to one in the CDB); and
- 5) sets the EN_BMS bit to one in the Background Control mode page in order to re-enable the background scan operation.

4.25 Association between commands and CbCS permission bits

Table 26 defines the CbCS permissions required for each command defined in this standard. The permissions shown in table 26 are defined in the PERMISSIONS BIT MASK field in the CbCS capability descriptor in a CbCS extension descriptor (see SPC-4). This standard does not define any permissions specific to block commands.

Table 26 – Associations between commands and CbCS permissions

Command	Permissions bit mask bits ^a				
	DATA READ	DATA WRITE	PARAM READ	PARAM WRITE	PHY ACC
COMPARE AND WRITE		1			
FORMAT UNIT		1		1	
GET LBA STATUS			1		
ORWRITE (16) / (32)		1			
POPULATE TOKEN	1				
PRE-FETCH (10) / (16)	1				
PREVENT ALLOW MEDIUM REMOVAL					1
READ (10) / (12) / (16) / (32)	1				
READ CAPACITY (10) / (16)			1		
READ DEFECT DATA (10) / (12)			1		
READ LONG (10) / (16)	1				
REASSIGN BLOCKS					1
RECEIVE ROD TOKEN INFORMATION	See SPC-4				
REPORT REFERRALS			1		
START STOP UNIT					1
SYNCHRONIZE CACHE (10) / (16)		1			
UNMAP		1			
VERIFY (10) / (12) / (16) / (32)	1				
WRITE (10) / (12) / (16) / (32)		1			
WRITE AND VERIFY (10) / (12) / (16) / (32)		1			
WRITE LONG (10) / (16)		1			
WRITE SAME (10) / (16) / (32)		1			
WRITE USING TOKEN		1			
XDWRITEREAD (10) / (32)	1	1			
XPWRITE (10) / (32)		1			

^a A device server shall only process a command shown in this table as specified by the CDB field of an extended CDB (see SPC-4) that contains a CbCS capability descriptor when all of the bits marked with a "1" in the row for that command are set to one in the PERMISSIONS BIT MASK field in that descriptor. The permissions bits represented by the empty cells in a row are ignored. If a device server receives a command specified by the CDB field of an extended CDB that does not contain the CbCS capability descriptor with all of the bits set to one as defined in this table, then the device server shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST.

4.26 Deferred microcode activation

After receiving a FORMAT UNIT command (see 5.3) or a START STOP UNIT command (see 5.25), a device server shall, prior to processing the command, activate any deferred microcode that has been downloaded as a result of a WRITE BUFFER command with the MODE field set to 0Eh (see SPC-4).

4.27 Model for uninterrupted sequences on LBA ranges

Direct access block devices may perform commands that require an uninterrupted sequence of actions to be performed on a specified range of LBAs. The uninterrupted sequence requirements are described in table 27. The uninterrupted sequences do not impact the processing of commands that access logical blocks other than those specified in the command requiring an uninterrupted sequence. The task attribute (see SAM-5) controls interactions between multiple commands. Commands with uninterrupted sequences on LBA ranges are shown in table 27.

Table 27 – Commands that require uninterrupted sequences

Command	Consistency enforcement	Reference
ORWRITE (16) ORWRITE (32)	The device server shall not perform any operations requested by any other command in the task set on logical blocks in the range specified by the command that requires an uninterrupted sequence of actions while performing the specified uninterrupted sequence of actions.	4.29
XDWRITEREAD (10) XDWRITEREAD (32)		5.47
XPWRITE (10) XPWRITE (32)		5.49
COMPARE AND WRITE	The device server shall not perform: <ul style="list-style-type: none"> a) any operations requested by any COMPARE AND WRITE command in the task set on logical blocks in the range specified by the command that requires an uninterrupted sequence of actions while performing the specified uninterrupted sequence of actions; b) any write operations to or unmap operations on logical blocks in the range specified by the command that requires an uninterrupted sequence of actions while performing the read operations specified in the uninterrupted sequence of actions; and c) any read operations or verify operations from logical blocks in the range specified by the command that requires an uninterrupted sequence of actions while performing the write operations specified in the uninterrupted sequence of actions. 	5.2

4.28 Referrals

4.28.1 Referrals overview

Referrals allow a logical unit to inform an application client that one or more user data segments (i.e., ranges of logical blocks) are accessible through target port group(s).

Support for referrals is indicated by the device server setting the R_SUP bit to one in the Extended INQUIRY Data VPD page (see SPC-4).

An application client may determine information on referrals by:

- a) issuing commands; or
- b) monitoring sense data returned as part of a completed command or a terminated command.

Figure 13 shows an example of how a logical unit informs an application client that one or more user data segments are accessible through target port groups.

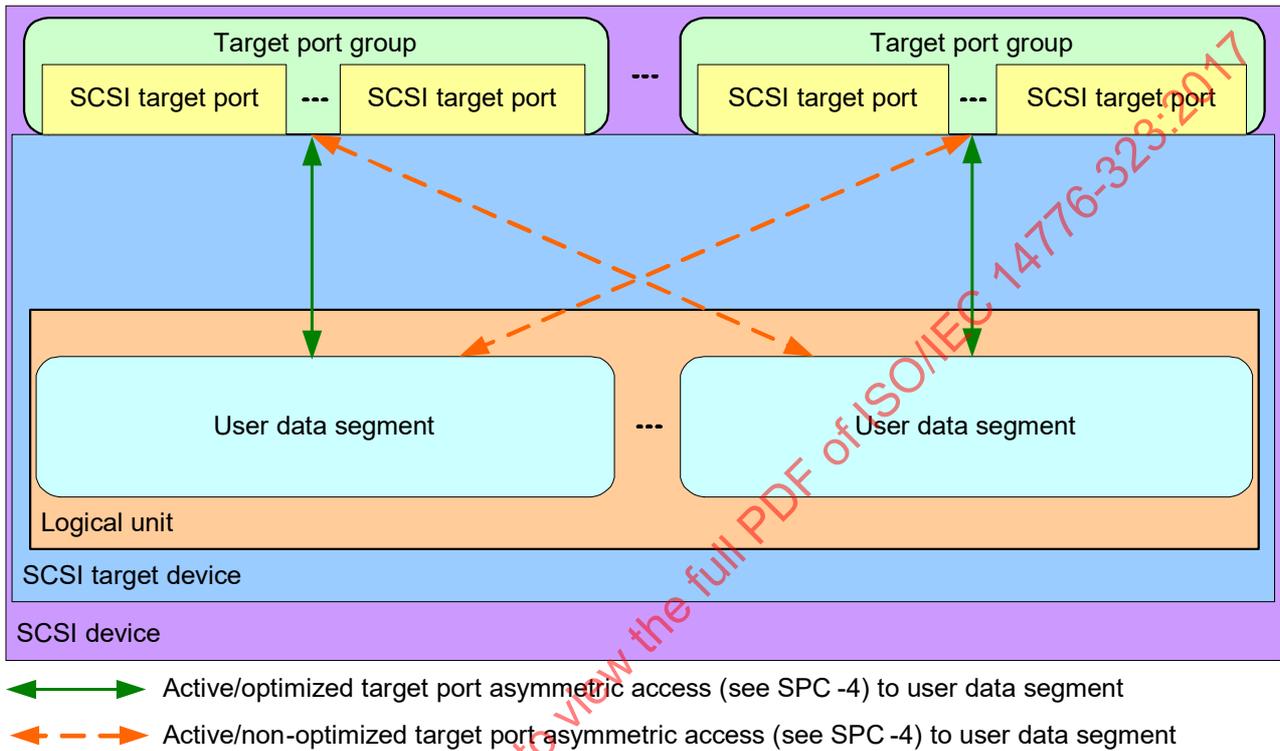


Figure 13 – Referrals

4.28.2 Discovering referrals

An application client may determine referrals information on a logical unit by:

- 1) determining if the R_SUP bit is set to one (i.e., the logical unit supports referrals) in the Extended INQUIRY Data VPD page (see SPC-4);
- 2) requesting the user data segment information from the Referrals VPD page (see 6.6.5);
- 3) requesting a list of target port groups by issuing a REPORT TARGET PORT GROUPS command (see SPC-4); and
- 4) either:
 - A) requesting referrals information by issuing a REPORT REFERRALS command (see 5.23); or
 - B) monitoring for referral information in sense data returned by the device server (see 4.28.3).

The following calculation is used to determine the first LBA for each user data segment within the range of LBAs indicated by the user data segment referral descriptors (see table 16) returned in:

- a) the REPORT REFERRALS parameter data (see table 85); or
- b) the user data segment referrals sense data descriptor (see 4.18.4):

$$\text{first LBA of the current user data segment} = \text{first LBA} + (\text{segment size} \times \text{segment multiplier})$$

where:

first LBA	the initial value is the first user data segment LBA specified in the user data segment referral descriptor (see table 16). Subsequent values, if any, are the first LBA of the previous user data segment;
segment size	the content of the USER DATA SEGMENT SIZE field (see 6.6.5); and
segment multiplier	the content of the USER DATA SEGMENT MULTIPLIER field (see 6.6.5).

If the content of the USER DATA SEGMENT SIZE field is greater than zero, and the content of the USER DATA SEGMENT MULTIPLIER field is greater than zero, then the following calculation may be used to determine the last LBA for each user data segment within the range of LBAs indicated by the user data segment referral descriptors (see table 16) returned in:

- a) the REPORT REFERRALS parameter data (see table 85); or
- b) the user data segment referrals sense data descriptor (see 4.18.4):

$$\text{last LBA of the current user data segment} = \text{first LBA} + (\text{segment size} - 1)$$

where:

first LBA	the first LBA of the current user data segment;
segment size	the content of the USER DATA SEGMENT SIZE field (see 6.6.5).

If the content of the USER DATA SEGMENT SIZE field is zero, then there is only one user data segment, and the last LBA of that user data segment is equal to the last LBA specified in the last USER DATA SEGMENT LBA field (see table 16).

See annex G for examples for discovering referrals.

4.28.3 Referrals in sense data

Returning referral information in sense data is enabled if the:

- a) R_SUP bit is set to one (i.e., the logical unit supports referrals) in the Extended INQUIRY Data VPD page (see SPC-4); and
- b) D_SENSE bit in the Control mode page is set to one (i.e., returning descriptor formatted sense data is enabled) (see SPC-4).

If reporting of referrals in sense data is enabled, a command completes without error, no other sense data is available within the logical unit, and the device server has an alternate I_T_L nexus that an application client should use to access at least one of the specified logical blocks, then the device server shall complete the command with GOOD status with the sense key set to COMPLETED, the additional sense code set to INSPECT REFERRALS SENSE DESCRIPTORS, and a user data segment referrals sense data descriptor (see 4.18.4).

The user data segment referral sense data descriptor (see 4.18.4) shall define the description of as many complete user data segments (i.e., one user data segment referral descriptor contains one complete user data segment) that fit in the maximum number of bytes allowed for sense data (i.e., 244 bytes or the maximum supported sense data length indicated in the Extended INQUIRY Data VPD page (see SPC-4)). If all the user data segments do not fit within the maximum number of bytes allowed for sense data, then:

- a) the device server shall set the NOT_ALL_R bit to one in the user data segment referral sense data descriptor (see 4.18.4); and
- b) the selection of which user data segments to include in the user data segment referral sense data descriptor is vendor specific.

Each user data segment referral sense data descriptor (see 4.18.4) contains information on alternate I_T_L nexuses to user data segments that the application client should use to access LBAs within the LBA range(s) indicated by the user data segments.

If reporting of referrals in sense data is enabled, the device server receives a command for which the device server is not able to access user data associated with the requested command, and the inaccessible user data is accessible through another target port group, then the device server shall terminate the command with CHECK CONDITION status with the sense key set to ABORTED COMMAND, the additional sense code set to INSPECT REFERRALS SENSE DESCRIPTORS, and a user data segment referral sense data descriptor.

The user data segment referral sense data descriptor shall, at a minimum, indicate the user data segment that contains the LBA of the first inaccessible logical block. Any other type of error that occurs while processing the command shall take precedence and be reported as described in this standard. If any other type of error occurs while the device server is processing the command, then processing that error shall take precedence over processing the command, and the device server shall report the error as described in this standard.

If reporting of referrals in sense data is disabled (see 4.28.1), the device server receives a command for which the device server is not able to access user data associated with the requested command, and the inaccessible user data is accessible through another target port group, then the device server shall terminate the command with CHECK CONDITION status with the sense key set to HARDWARE ERROR and the additional sense code set to INTERNAL TARGET FAILURE.

4.29 ORWRITE commands

4.29.1 ORWRITE commands overview

The ORWRITE commands (see 5.5 and 5.6) provide a mechanism for an application client to manipulate bitmap structures on direct access block devices.

An ORWRITE command shall be processed by the device server performing the following as an uninterrupted sequence of actions (see 4.27):

- 1) perform read operations from the LBAs specified by this command;
- 2) transfer the specified number of logical blocks from the Data-Out Buffer;
- 3) perform the specified Boolean arithmetic function on:
 - A) the user data contained in the logical blocks from read operations; and
 - B) the user data contained in the logical blocks transferred from the Data-Out Buffer;
- 4) store the results of the Boolean arithmetic function in a bitmap buffer;
- 5) generate new protection information, if any, from the stored results;
- 6) store the generated protection information, if any, into the bitmap buffer; and
- 7) perform write operations using the updated logical block data from the bitmap buffer.

If the check of the protection information from the read operations is successful (see table 46), and the check of the protection information transferred from the Data-Out Buffer is successful (see table 47), then the device server shall generate the new protection information (see 4.22) as follows:

- a) set the LOGICAL BLOCK GUARD field to the CRC (see 4.22.4) generated from the bitmap buffer by the device server;
- b) set the LOGICAL BLOCK REFERENCE TAG field to the LOGICAL BLOCK REFERENCE TAG field received from the Data-Out Buffer; and
- c) set the LOGICAL BLOCK APPLICATION TAG field to the LOGICAL BLOCK APPLICATION TAG field received from the Data-Out Buffer.

In order to support the manipulation of bitmap structures:

- a) the ORWRITE (16) command supports the set operation (see 4.29.4); and
- b) the ORWRITE (32) command supports:
 - A) the set operation; and
 - B) the change generation and clear operation (see 4.29.3).

4.29.2 ORWgeneration code

4.29.2.1 ORWgeneration code overview

The ORWRITE commands use a generation code for synchronization. The device server shall establish and maintain the following generation codes:

- a) a current ORWgeneration code; and
- b) a previous ORWgeneration code.

Subsequent ORWRITE command processing by the device server is dependent on comparisons involving the ORWgeneration codes. Changes in these ORWgeneration codes define a synchronization point in the management of the bitmap.

4.29.2.2 ORWgeneration code processing

The device server shall maintain at least one current ORWRITE processing policy. The device server may support more than one ORWRITE processing policy (see table 28 in 4.29.4).

When processing an ORWRITE (32) command (see 5.6), the device server compares the value in the EXPECTED ORWGENERATION field in the CDB to the current ORWgeneration code (see 4.29.2.1), and:

- a) if the two values are equal, then the device server continues processing the ORWRITE (32) command as described in table 28 for the set operation and as described in 4.29.3 for the change generation and clear operation; or
- b) if the two values are not equal, then:
 - A) for a set operation, the current ORWRITE processing policy (see table 28) determines how the device server continues processing the ORWRITE (32) command; and
 - B) for a change generation and clear operation, the device server terminates the ORWRITE (32) command (see 4.29.3).

If the device server supports both the ORWRITE (16) command (see 5.5) and the ORWRITE (32) command, then the device server shall process all ORWRITE (16) commands as if they contained an EXPECTED ORWGENERATION field set to zero.

4.29.3 Change generation and clear operation

The change generation portion of the change generation and clear operation is used to establish a point of synchronization. The clear portion of the change generation and clear operation is used to set zero or more bits in the bitmap structure to zero.

The device server performs a change generation and clear operation if:

- a) the BMOP field in the ORWRITE (32) command (see 5.6) is set to 001b; and
- b) the value in the EXPECTED ORWGENERATION field is equal to the current ORWgeneration code in the device server.

If the device server performs a change generation and clear operation, then the device server shall perform the following as an uninterrupted sequence:

- 1) perform read operations from the LBAs specified by this command;
- 2) transfer the specified logical blocks from the Data-Out Buffer;
- 3) perform an AND operation (see 3.1.3) on the user data contained in the logical blocks from the read operations and the user data contained in the logical blocks transferred from the Data-Out Buffer;
- 4) store the results of the AND operation in a bitmap buffer;
- 5) generate new protection information, if any, from the stored results;
- 6) store the generated protection information, if any, into the bitmap buffer;
- 7) perform write operations using the updated logical block data from the bitmap buffer;
- 8) set the current ORWRITE processing policy to the value in the PREVIOUS GENERATION PROCESSING field in the ORWRITE (32) command;
- 9) set the previous ORWgeneration code (see 4.29.2) to the current ORWgeneration code in the device server; and
- 10) set the current ORWgeneration code (see 4.29.2) to the value in the NEW ORWGENERATION field in the ORWRITE (32) command.

If the value in the EXPECTED ORWGENERATION field is not equal to the current ORWgeneration code, then the device server shall terminate the ORWRITE (32) command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to ORWRITE GENERATION DOES NOT MATCH.

If a power on or hard reset condition occurs, then the device server shall set:

- a) the current ORWgeneration code to zero;

- b) the previous ORWgeneration code to zero; and
- c) the current ORWRITE processing policy to 7h.

The device server shall preserve the following across a logical unit reset:

- a) the current ORWgeneration code;
- b) the previous ORWgeneration code; and
- c) the current ORWRITE processing policy.

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC 14776-323:2017

4.29.4 Set operation

The set operation is used to set zero or more bits in the bitmap structure to one.

The device server performs a set operation for an ORWRITE command (see 5.5 and 5.6) if the BMOP field in the CDB is set to 000b.

The device server shall perform a set operation by performing the actions specified in table 28, which shows the current ORWgeneration code, the previous ORWgeneration code, and the device server's current ORWRITE processing policy (see 4.29.3).

Table 28 – Performing an ORWRITE set operation

Current ORWRITE processing policy	The value in the EXPECTED ORWGENERATION field matches		
	Current ORWgeneration code	Previous ORWgeneration code	Any other value
0h	PA	PA	CCG
1h	Reserved		
2h	PA	DN	CCG
3h	PA	PA	PA
4h	Reserved		
5h	PA	DN	DN
6h	Reserved		
7h	PA	CCG	CCG
8h to Fh	Reserved		

Key:

PA = the device server shall perform the following as an uninterrupted sequence:

- 1) perform read operations from the LBAs specified by the command;
- 2) transfer the specified logical blocks from the Data-Out Buffer;
- 3) perform an OR operation on the logical blocks from the read operations and the user data contained in the logical blocks transferred from the Data-Out Buffer;
- 4) store the results of the OR operation in a bitmap buffer;
- 5) generate new protection information, if any, from the stored results;
- 6) store the generated protection information, if any, into the bitmap buffer; and
- 7) perform write operations using the updated logical block data (i.e., those containing the user data resulting from the OR operation and the generated protection information, if any) from the bitmap buffer.

DN = the device server shall discard the contents of the Data-Out Buffer and shall complete the command with GOOD status.

CCG = the device server shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to ORWRITE GENERATION DOES NOT MATCH

4.30 Block device ROD token operations

4.30.1 Block device ROD token operations overview

Application clients request that block device ROD token operations (see SPC-4) be performed using the commands summarized in this subclause or the commands specified in SPC-4.

Copy managers (see SPC-4) that implement the POPULATE TOKEN command (see 5.7) or the WRITE USING TOKEN command (see 5.46) shall implement the following:

- a) the POPULATE TOKEN command;
- b) the WRITE USING TOKEN command;
- c) the RECEIVE COPY STATUS (LID4) command (see SPC-4);
- d) the RECEIVE ROD TOKEN INFORMATION command (see SPC-4 and 5.22); and
- e) the Third-party Copy VPD page (see 6.6.6) containing at least one Block Device ROD Token Limits descriptor (see 6.6.6.3).

The POPULATE TOKEN command may cause the copy manager to create zero or one point in time ROD tokens. If the POPULATE TOKEN command causes one point in time ROD token to be created, then this point in time ROD token may be retrieved by an application client using the RECEIVE ROD TOKEN INFORMATION command.

The WRITE USING TOKEN command causes the copy manager to transfer the data represented by the specified ROD token (i.e., the data represented by the ROD token retrieved using the RECEIVE ROD TOKEN INFORMATION command or the data represented by the block device zero ROD token).

The copy manager manages the point in time ROD token.

After the copy manager begins processing a POPULATE TOKEN command or a WRITE USING TOKEN command, the copy manager shall preserve information for return in response to a RECEIVE ROD TOKEN INFORMATION command as defined in SPC-4.

Block device range descriptor lists (see 5.7.3) contain non-overlapping block device range descriptors and are used by the application client to specify:

- a) the logical blocks to include in the data represented by the ROD token;
- b) the sequence of the logical blocks in the data represented by the ROD token (e.g., the first logical block represented by the LBA described in the first block device range descriptor is placed at the beginning of the data represented by the ROD token, and the first logical block represented by the LBA described in the second block device range descriptor is placed in the data represented by the ROD token immediately following the last logical block represented by the LBA described in the first block device range descriptor);
- c) the logical blocks to be written from the data represented by the ROD token; and
- d) the sequence of the logical blocks written from the data represented by the ROD token (e.g., the first logical block represented by the LBA described in the first block device range descriptor is written from the beginning of the data represented by the ROD token, and the first logical block represented by the LBA described in the second block device range descriptor is written from data represented by the ROD token immediately following the data written to the last logical block represented by the LBA described in the first block device range descriptor).

If the copy manager uses out of order transfers to create the representation of data for the ROD token, then the TRANSFER COUNT field in the parameter data returned in the response to a RECEIVE ROD TOKEN INFORMATION command with a list identifier that specifies a POPULATE TOKEN command (see 5.22.2) shall be based only on the contiguous transfers that complete without error starting at the first LBA specified by the first block device range descriptor (i.e., any transfers completed without error beyond the first incomplete or unsuccessful transfer shall not contribute to the computation of the value in the TRANSFER COUNT field).

If the copy manager uses out of order transfers to write from the data represented by the ROD token, then the TRANSFER COUNT field in the parameter data returned in response to a RECEIVE ROD TOKEN INFORMATION command with a list identifier that specifies a WRITE USING TOKEN command (see 5.22.3) shall be based only on the contiguous transfers that complete without error starting at the first LBA specified

by the first block device range descriptor (i.e., any transfers completed without error beyond the first incomplete or unsuccessful transfer shall not contribute to the computation of the value in the TRANSFER COUNT field).

4.30.2 POPULATE TOKEN command and WRITE USING TOKEN command completion

As part of completing a block device token operation originated by a POPULATE TOKEN command (see 5.7) or a WRITE USING TOKEN command (see 5.46), the copy manager shall compute the residual by subtracting the sum of the contents of the NUMBER OF LOGICAL BLOCK fields in all of the complete block device range descriptors of the parameter list (see 5.7.3) from the TRANSFER COUNT field in the parameter data returned in response to a RECEIVE ROD TOKEN INFORMATION command (see 5.22.2 and 5.22.3).

If the POPULATE TOKEN command was received with the IMMED bit set to zero, and the residual is negative, then the copy manager shall:

- a) terminate the command with CHECK CONDITION status, with the additional sense code set to COPY TARGET DEVICE DATA UNDERRUN, the sense key set to:
 - A) COPY ABORTED, if the transfer count is not zero; or
 - B) ILLEGAL REQUEST, if the transfer count is zero,and report the transfer count in the INFORMATION field (see SPC-4); or
- b) complete the command with GOOD status and return parameter data for the RECEIVE ROD TOKEN INFORMATION command received on the same I_T nexus with a matching LIST IDENTIFIER field with:
 - A) the COPY OPERATION STATUS field set to 03h or 60h (see SPC-4);
 - B) the EXTENDED COPY COMPLETION STATUS field set to CHECK CONDITION (see SAM-5); and
 - C) the SENSE DATA field with the additional sense code set to COPY TARGET DEVICE DATA UNDERRUN, the sense key set to:
 - a) COPY ABORTED, if the transfer count is not zero; or
 - b) ILLEGAL REQUEST, if the transfer count is zero;and report the transfer count in the INFORMATION field (see SPC-4).

If the WRITE USING TOKEN command was received with the IMMED bit set to zero, and the residual is negative, then the copy manager shall terminate the command with CHECK CONDITION status, the additional sense code set to COPY TARGET DEVICE DATA UNDERRUN, the sense key set to:

- a) COPY ABORTED, if the transfer count is not zero; or
- b) ILLEGAL REQUEST, if the transfer count is zero,

and report the transfer count in the INFORMATION field (see SPC-4).

If the POPULATE TOKEN command or WRITE USING TOKEN command was received with the IMMED bit set to one, and the residual is negative, then the copy manager shall return parameter data for the RECEIVE ROD TOKEN INFORMATION command received on the same I_T nexus with a matching LIST IDENTIFIER field with:

- a) the COPY OPERATION STATUS field set to 03h or 60h (see SPC-4);
- b) the EXTENDED COPY COMPLETION STATUS field set to CHECK CONDITION status (see SAM-5); and
- c) the SENSE DATA field with the additional sense code set to COPY TARGET DEVICE DATA UNDERRUN, the sense key set to:
 - A) COPY ABORTED, if the transfer count is not zero; or
 - B) ILLEGAL REQUEST, if the transfer count is zero,and report the transfer count in the INFORMATION field (see SPC-4).

4.30.3 Block device specific ROD tokens

Block device specific ROD token types (see SPC-4) are shown in table 29.

Table 29 – ROD token type values

ROD token type	Description	Reference
FF00_0000h to FFFF_0000h	Reserved	
FFFF_0001h	Block device zero ROD token	4.30.4
FFFF_0002h to FFFF_FFEFh	Reserved	

4.30.4 Block device zero ROD token

The block device zero ROD token represents user data in which all bits are set to zero and protection information, if any, is set to FFFF_FFFF_FFFF_FFFFh. If user data with all bits set to zero and protection information, if any, set to FFFF_FFFF_FFFF_FFFFh, is represented by a ROD token, then, in response to a RECEIVE ROD TOKEN INFORMATION command in which the LIST IDENTIFIER field specifies a POPULATE TOKEN command, the copy manager may or may not return a ROD token that is the block device zero ROD token. The block device zero ROD token format is shown in table 30.

Table 30 – Block device zero ROD token format

Byte	Bit	7	6	5	4	3	2	1	0
0	(MSB)	ROD TOKEN TYPE (FFFF_0001h)							
...									
3									
4		Reserved							
5		ROD TOKEN LENGTH (01F8h)							
6	(MSB)								
7									
8		Reserved							
...									
511									

The logical block length associated with the block device zero ROD token is that of the direct access block device to which the data is being written (e.g., if a block device zero ROD token is used to write to a logical unit that has 642-byte logical blocks, then the logical block length of the block device zero ROD token is 642 bytes).

The ROD TOKEN TYPE field is defined in SPC-4 and shall be set to the value shown in table 30 for the block device zero ROD token.

The ROD TOKEN LENGTH field is defined in SPC-4 and shall be set to the value shown in table 30 for the block device zero ROD token.

4.30.5 ROD token device type specific data

The device type specific data for ROD tokens (see SPC-4) created by a copy manager for a direct access block devices (see 6.7):

- a) provides information about the logical unit at the time that the ROD token was created; and
- b) is a subset of the parameter data returned by the READ CAPACITY (16) command (see 5.16) for the logical unit that contains the copy manager that created the ROD token.

If the READ CAPACITY (16) parameter data changes so that the copy manager that created the ROD token is no longer able to access the data represented by the ROD token, then that copy manager shall invalidate the ROD token.

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC 14776-323:2017

5 Commands for direct access block devices

5.1 Commands for direct access block devices overview

The commands for direct access block devices are listed in table 31.

Table 31 – Commands for direct access block devices (part 1 of 4)

Command	Operation code ^a	Type	LBACT	Reference
ACCESS CONTROL IN	86h	O	n/a	SPC-4
ACCESS CONTROL OUT	87h	O	n/a	SPC-4
ATA PASS-THROUGH (12)	A1h	O	n/a	SAT-3
ATA PASS-THROUGH (16)	85h	O	n/a	SAT-3
CHANGE ALIASES	A4h/0Bh	O	n/a	SPC-4
COMPARE AND WRITE	89h	O	R, W	5.2
EXTENDED COPY	83h	O	n/a	SPC-4
FORMAT UNIT	04h	M	Z	5.3
GET LBA STATUS	9Eh/12h	O	n/a	5.4
INQUIRY	12h	M	n/a	SPC-4
LOG SELECT	4Ch	O	n/a	SPC-4
LOG SENSE	4Dh	O	n/a	SPC-4
MAINTENANCE IN	A3h/00h to 04h A3h/06h to 09h	X	n/a	SPC-4 SCC-2
MAINTENANCE OUT	A4h/00h to 05h A4h/07h to 09h	X	n/a	SPC-4 SCC-2
MODE SELECT (6)	15h	O	n/a	SPC-4
MODE SELECT (10)	55h	O	n/a	SPC-4
MODE SENSE (6)	1Ah	O	n/a	SPC-4
MODE SENSE (10)	5Ah	O	n/a	SPC-4
ORWRITE (16)	8Bh	O	R, W	5.5
ORWRITE (32)	7Fh/000Eh	O	R, W	5.6
PERSISTENT RESERVE IN	5Eh	O	n/a	SPC-4
PERSISTENT RESERVE OUT	5Fh	O	n/a	SPC-4
POPULATE TOKEN	83h/10h	O	n/a	5.7
PRE-FETCH (10)	34h	O	R	5.8

Key:

O	= optional	V	= verify command
M	= mandatory	W	= write command
X	= implementation requirements are defined in the reference	Z	= other command
R	= read command	PI	= protection information
U	= unmap command	LBACT	= logical block access command type (see 4.2.2)

^a If a command is defined by a combination of operation code and service action, then the operation code value is shown preceding a slash and the service action value is shown after the slash.

Table 31 – Commands for direct access block devices (part 2 of 4)

Command	Operation code ^a	Type	LBACT	Reference
PRE-FETCH (16)	90h	O	R	5.9
PREVENT ALLOW MEDIUM REMOVAL	1Eh	O	n/a	5.10
READ (10)	28h	M	R	5.11
READ (12)	A8h	O	R	5.12
READ (16)	88h	M	R	5.13
READ (32)	7Fh/0009h	O	R	5.14
READ ATTRIBUTE	8Ch	O	n/a	SPC-4
READ BUFFER	3Ch	O	n/a	SPC-4
READ CAPACITY (10)	25h	M	n/a	5.15
READ CAPACITY (16)	9Eh/10h	M	n/a	5.16
READ DEFECT DATA (10)	37h	O	n/a	5.17
READ DEFECT DATA (12)	B7h	O	n/a	5.18
READ LONG (10)	3Eh	O	Z	5.19
READ LONG (16)	9Eh/11h	O	Z	5.20
REASSIGN BLOCKS	07h	O	Z	5.21
RECEIVE COPY RESULTS	84h	O	n/a	SPC-4
RECEIVE DIAGNOSTIC RESULTS	1Ch	X	n/a	SPC-4 6.3 4.30
RECEIVE ROD TOKEN INFORMATION	84h/07h	X	n/a	SPC-4 5.22
REDUNDANCY GROUP IN	BAh	X	n/a	SCC-2
REDUNDANCY GROUP OUT	BBh	X	n/a	SCC-2
REMOVE I_T NEXUS	A4h/0Ch	O	n/a	SPC-4
REPORT REFERRALS	9Eh/13h	O	n/a	5.23
REPORT ALIASES	A3h/0Bh	O	n/a	SPC-4
REPORT IDENTIFYING INFORMATION	A3h/05h	O	n/a	SPC-4
REPORT LUNS	A0h	M	n/a	SPC-4
REPORT PRIORITY	A3h/0Eh	O	n/a	SPC-4
REPORT SUPPORTED OPERATION CODES	A3h/0Ch	O	n/a	SPC-4
REPORT SUPPORTED TASK MANAGEMENT FUNCTIONS	A3h/0Dh	O	n/a	SPC-4
REPORT TARGET PORT GROUPS	A3h/0Ah	O	n/a	SPC-4
Key:				
O = optional		V = verify command		
M = mandatory		W = write command		
X = implementation requirements are defined in the reference		Z = other command		
R = read command		PI = protection information		
U = unmap command		LBACT= logical block access command type (see 4.2.2)		
^a If a command is defined by a combination of operation code and service action, then the operation code value is shown preceding a slash and the service action value is shown after the slash.				

Table 31 – Commands for direct access block devices (part 3 of 4)

Command	Operation code ^a	Type	LBACT	Reference
REQUEST SENSE	03h	M	n/a	SPC-4
SANITIZE	48h	O	Z	5.24
SECURITY PROTOCOL IN	A2h	O	n/a	SPC-4
SECURITY PROTOCOL OUT	B5h	O	n/a	SPC-4
SEND DIAGNOSTIC	1Dh	O	n/a	SPC-4
SET IDENTIFYING INFORMATION	A4h/06h	O	n/a	SPC-4
SET PRIORITY	A4h/0Eh	O	n/a	SPC-4
SET TARGET PORT GROUPS	A4h/0Ah	O	n/a	SPC-4
SPARE IN	BCh	X	n/a	SCC-2
SPARE OUT	BDh	X	n/a	SCC-2
START STOP UNIT	1Bh	O	n/a	5.25
SYNCHRONIZE CACHE (10)	35h	O	W	5.26
SYNCHRONIZE CACHE (16)	91h	O	W	5.27
TEST UNIT READY	00h	M	n/a	SPC-4
UNMAP	42h	X	U	5.28 4.7
VERIFY (10)	2Fh	O	V, W	5.29
VERIFY (12)	AFh	O	V, W	5.30
VERIFY (16)	8Fh	O	V, W	5.31
VERIFY (32)	7Fh/000Ah	O	V, W	5.32
VOLUME SET IN	BEh	X	n/a	SCC-2
VOLUME SET OUT	BFh	X	n/a	SCC-2
WRITE (10)	2Ah	O	W	5.33
WRITE (12)	AAh	O	W	5.34
WRITE (16)	8Ah	O	W	5.35
WRITE (32)	7Fh/000Bh	O	W	5.36
WRITE AND VERIFY (10)	2Eh	O	V, W	5.37
WRITE AND VERIFY (12)	A Eh	O	V, W	5.38
WRITE AND VERIFY (16)	8 Eh	O	V, W	5.39
WRITE AND VERIFY (32)	7Fh/000Ch	O	V, W	5.40
WRITE ATTRIBUTE	8Dh	O	n/a	SPC-4
WRITE BUFFER	3Bh	O	n/a	SPC-4
Key:				
O = optional		V = verify command		
M = mandatory		W = write command		
X = implementation requirements are defined in the reference		Z = other command		
R = read command		PI = protection information		
U = unmap command		LBACT= logical block access command type (see 4.2.2)		
^a If a command is defined by a combination of operation code and service action, then the operation code value is shown preceding a slash and the service action value is shown after the slash.				

Table 31 – Commands for direct access block devices (part 4 of 4)

Command	Operation code ^a	Type	LBACT	Reference										
WRITE LONG (10)	3Fh	O	Z	5.41										
WRITE LONG (16)	9Fh/11h	O	Z	5.42										
WRITE SAME (10)	41h	X	U, W	5.43 4.7										
WRITE SAME (16)	93h	X	U, W	5.44 4.7										
WRITE SAME (32)	7Fh/000Dh	X	U, W	5.45 4.7										
WRITE USING TOKEN	83h/11h	X	Z	5.46 4.30										
XDWRITEREAD (10)	53h	O	R, W	5.47										
XDWRITEREAD (32)	7Fh/0007h	O	R, W	5.48										
XPWRITE (10)	51h	O	R, W	5.49										
XPWRITE (32)	7Fh/0006h	O	R, W	5.50										
<p>Note 1 - The following operation codes are obsolete: 01h, 08h, 0Ah, 0Bh, 16h, 17h, 18h, 2Bh, 30h, 31h, 32h, 33h, 36h, 39h, 3Ah, 40h, 50h, 52h, 56h, 57h, 80h, 81h, 82h, 92h, A7h, B3h, and B4h.</p> <p>Note 2 - The following operation codes are vendor specific: 02h, 05h, 06h, 09h, 0Ch, 0Dh, 0Eh, 0Fh, 10h, 11h, 13h, 14h, 19h, 20h, 21h, 22h, 23h, 24h, 26h, 27h, 29h, 2Ch, 2Dh, and C0h to FFh.</p> <p>Note 3 - A complete summary of operation codes is available at http://www.t10.org/lists/2op.htm. The summary includes information about obsolete commands.</p>														
<p>Key:</p> <table style="width: 100%; border: none;"> <tr> <td style="width: 50%;">O = optional</td> <td style="width: 50%;">V = verify command</td> </tr> <tr> <td>M = mandatory</td> <td>W = write command</td> </tr> <tr> <td>X = implementation requirements are defined in the reference</td> <td>Z = other command</td> </tr> <tr> <td>R = read command</td> <td>PI = protection information</td> </tr> <tr> <td>U = unmap command</td> <td>LBACT= logical block access command type (see 4.2.2)</td> </tr> </table>					O = optional	V = verify command	M = mandatory	W = write command	X = implementation requirements are defined in the reference	Z = other command	R = read command	PI = protection information	U = unmap command	LBACT= logical block access command type (see 4.2.2)
O = optional	V = verify command													
M = mandatory	W = write command													
X = implementation requirements are defined in the reference	Z = other command													
R = read command	PI = protection information													
U = unmap command	LBACT= logical block access command type (see 4.2.2)													
<p>^a If a command is defined by a combination of operation code and service action, then the operation code value is shown preceding a slash and the service action value is shown after the slash.</p>														

5.2 COMPARE AND WRITE command

The COMPARE AND WRITE command (see table 32) requests that the device server perform the following as an uninterrupted sequence of actions (see 4.27):

- 1) perform read operations from the specified LBAs;
- 2) perform a compare operation on only the user data (i.e., not on the protection information) from:
 - A) the read operations; and
 - B) the compare instance transferred from the Data-Out Buffer;
- 3) if the compare operation indicates a match, then perform the following operations:
 - 1) check the protection information, if any, in the write instance transferred from the Data-Out Buffer based on the contents of the WRPROTECT field as defined in table 107; and
 - 2) perform write operations to the LBAs specified by this command using the write instance;

and
- 4) if the compare operation does not indicate a match, then terminate the command with CHECK CONDITION status with the sense key set to MISCOMPARE and the additional sense code set to

MISCOMPARE DURING VERIFY OPERATION. In the sense data (see 4.18 and SPC-4) the offset from the start of the Data-Out Buffer to the first byte of data that was not equal shall be reported in the INFORMATION field.

The Data-Out Buffer contains two instances of logical block data:

- 1) the compare instance, in which:
 - A) the user data is used for the compare operation; and
 - B) the protection information, if any, is not used;
 and
- 2) the write instance, in which:
 - A) the user data is used for the write operations; and
 - B) the protection information, if any, is used for the write operations.

Table 32 – COMPARE AND WRITE command

Bit Byte	7	6	5	4	3	2	1	0
0	OPERATION CODE (89h)							
1	WRPROTECT			DPO	FUA	Reserved		
2	(MSB)							
...	LOGICAL BLOCK ADDRESS							
9	(LSB)							
10	Reserved							
...	Reserved							
12	Reserved							
13	NUMBER OF LOGICAL BLOCKS							
14	Reserved			GROUP NUMBER				
15	CONTROL							

The OPERATION CODE field is defined in SPC-4 and shall be set to the value shown in table 32 for the COMPARE AND WRITE command.

See the WRITE (10) command for the definition of the WRPROTECT field.

See the READ (10) command (see 5.11) for the definition of the DPO bit. See the READ (10) command (see 5.11) for the definition of the FUA bit specifying behavior for the read operations. See the WRITE (10) command (see 5.33) for the definition of the FUA bit specifying behavior for the write operations.

The LOGICAL BLOCK ADDRESS field specifies the LBA of the first logical block (see 4.5) to be accessed by the device server (e.g., the first LBA accessed by both a read operation and a write operation). If the specified LBA exceeds the capacity of the medium, then the device server shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to LOGICAL BLOCK ADDRESS OUT OF RANGE.

The NUMBER OF LOGICAL BLOCKS field specifies:

- a) the number of contiguous logical blocks on which read operations shall be performed, starting with the LBA specified by the LOGICAL BLOCK ADDRESS field;
- b) the number of contiguous logical blocks that shall be transferred from the Data-Out Buffer for the compare operation; and
- c) if the compare operation indicates a match, then the number of contiguous logical blocks that shall be transferred from the Data-Out Buffer and on which write operations shall be performed, starting with the LBA specified by the LOGICAL BLOCK ADDRESS field.

A NUMBER OF LOGICAL BLOCKS field set to zero specifies that no read operations shall be performed, no logical block data shall be transferred from the Data-Out Buffer, no compare operations shall be performed, and no write operations shall be performed. This condition shall not be considered an error.

If the specified LBA and the specified number of logical blocks exceed the capacity of the medium (see 4.5), then the device server shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to LOGICAL BLOCK ADDRESS OUT OF RANGE.

If the number of logical blocks exceeds the value in the MAXIMUM COMPARE AND WRITE LENGTH field in the Block Limits VPD page (see 6.6.3), then the device server shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

See the PRE-FETCH (10) command (see 5.8) and 4.23 for the definition of the GROUP NUMBER field.

The CONTROL byte is defined in SAM-5.

5.3 FORMAT UNIT command

5.3.1 FORMAT UNIT command overview

The FORMAT UNIT command (see table 33) requests that the device server format the medium into application client accessible logical blocks as specified in the number of logical blocks and logical block length values received in the last mode parameter block descriptor (see 6.5.2) in a MODE SELECT command (see SPC-4). In addition, the device server may certify the medium and create control structures for the management of the medium and defects. The degree that the medium is altered by this command is vendor specific.

If a device server receives a FORMAT UNIT command before receiving a MODE SELECT command with a mode parameter block descriptor, then the device server shall use the number of logical blocks and logical block length at which the logical unit is currently formatted (i.e., no change is made to the number of logical blocks and the logical block length of the logical unit during the format operation).

The device server shall handle any deferred microcode as specified in 4.26.

Before performing the operation specified by this command, the device server shall stop all:

- a) enabled power condition timers (see SPC-4);
- b) timers for enabled background scan operations (see 4.24); and
- c) timers or counters enabled for device-specific background functions.

After the operation is complete, the device server shall reinitialize and restart all enabled timers and counters for power conditions and background functions.

While performing a format operation, the device server shall:

- a) process commands already in a task set when a FORMAT UNIT command is received in a vendor specific manner;
- b) process an INQUIRY command by returning parameter data based on the condition of the logical unit before beginning the FORMAT UNIT command (i.e., INQUIRY data shall not change until after successful completion of a format operation);
- c) process a REQUEST SENSE command by returning parameter data containing sense data with the sense key set to NOT READY, the additional sense code set to LOGICAL UNIT NOT READY, FORMAT IN PROGRESS, and the PROGRESS INDICATION field in the sense data (see SPC-4) set to indicate the progress of the format operation;
- d) process REPORT LUNS commands; and
- e) terminate all commands, except INQUIRY commands, REPORT LUNS commands, and REQUEST SENSE commands, with CHECK CONDITION status with the sense key set to NOT READY and the additional sense code set to LOGICAL UNIT NOT READY, FORMAT IN PROGRESS.

NOTE 5 - The progress indication is available to the application client during a format operation as a means to determine the logical unit's format progress.

The simplest form of the FORMAT UNIT command (i.e., a FORMAT UNIT command with no parameter data) accomplishes medium formatting with little application client control over medium defect management. The device server implementation determines the degree of medium defect management that is to be performed. Additional forms of this command increase the application client's control over medium defect management (see 4.13). The application client may specify:

- a) that the device server clear the existing GLIST;
- b) a list of address descriptors that the device server adds to the GLIST;
- c) that the device server enable a certification operation that adds address descriptors for physical blocks with medium defects discovered during the certification operation to the GLIST; and
- d) the behavior of the device server if it is not able to access the PLIST or GLIST or determine whether one or both of them exists.

Following a successful format operation:

- a) if the logical unit is fully provisioned (i.e., the LBPME bit in the READ CAPACITY (16) parameter data is set to zero), then all LBAs in the logical unit are mapped (see 4.7.2); or
- b) if the logical unit supports logical block provisioning management (i.e., the LBPME bit is set to one), then:
 - A) if the LBPRZ bit in the READ CAPACITY (16) parameter data is set to zero, then each LBA in the logical unit shall be either:
 - a) mapped, if an initialization pattern was specified that does not match the vendor specific data returned by a read command for an unmapped LBA (see 4.7.4.5); or
 - b) unmapped (see 4.7.3.2 and 4.7.3.3), if no initialization pattern was specified or an initialization pattern was specified that matches the vendor specific data returned by a read command for an unmapped LBA (see 4.7.4.5); and
 - B) if the LBPRZ bit is set to one, then each LBA in the logical unit:
 - a) shall be mapped, if the format operation did not initialize the user data to all zeroes for the logical block referenced by that LBA;
 - b) shall be unmapped (see 4.7.3.2 and 4.7.3.3), if the format operation initialized the user data to all zeroes for the logical blocks referenced by all valid LBAs in the logical unit; or
 - c) may be unmapped (see 4.7.3.2 and 4.7.3.3), if the format operation initialized the user data to all zeroes for the logical block referenced by that LBA, and the format operation did not initialize the user data to all zeroes for the logical blocks referenced by all valid LBAs in the logical unit.

Table 33 defines the FORMAT UNIT command.

Table 33 – FORMAT UNIT command

Bit Byte	7	6	5	4	3	2	1	0
0	OPERATION CODE (04h)							
1	FMTPINFO		LONGLIST	FMTDATA	CMPLST	DEFECT LIST FORMAT		
2	Vendor specific							
3	Obsolete							
4								
5	CONTROL							

The OPERATION CODE field is defined in SPC-4 and shall be set to the value shown in table 33 for the FORMAT UNIT command.

The combination (see table 38) of the format protection information (FMTPINFO) field and the PROTECTION FIELD USAGE field in the parameter list header (see 5.3.2.2) specifies whether or not the device server enables or disables the use of protection information.

A **LONGLIST** bit set to zero specifies that the parameter list, if any, contains a short parameter list header as defined in table 36. A **LONGLIST** bit set to one specifies that the parameter list, if any, contains a long parameter list header as defined in table 37. If the **FMTDATA** bit is set to zero, then the **LONGLIST** bit shall be ignored.

A format data (**FMTDATA**) bit set to one specifies that the **FORMAT UNIT** parameter list (see 5.3.2) shall be transferred from the Data-Out Buffer. A **FMTDATA** bit set to zero specifies that no parameter list be transferred from the Data-Out Buffer. If the **FMTDATA** bit is set to zero and the **FMTPINFO** field is not set to zero, then the device server shall terminate the command with **CHECK CONDITION** status with the sense key set to **ILLEGAL REQUEST** and the additional sense code set to **INVALID FIELD IN CDB**.

Following a successful format operation, the **PROT_EN** bit and the **P_TYPE** field in the **READ CAPACITY (16)** parameter data (see 5.16.2) indicate the type of protection currently in effect on the logical unit.

If protection information is written during a format operation (i.e., the **FMTPINFO** field is set to a value greater than zero), then protection information shall be written to a default value of **FFFF_FFFF_FFFF_FFFFh**.

A complete list (**CMPLST**) bit set to zero specifies that the device server shall add the defect list included in the **FORMAT UNIT** parameter list to the existing **GLIST** (see 4.13). A **CMPLST** bit set to one specifies that the device server shall replace the existing **GLIST** with the defect list, if any, included in the **FORMAT UNIT** parameter list.

If the **FMTDATA** bit is set to zero, then the **CMPLST** bit shall be ignored.

If the **FMTDATA** bit is set to one, then the **DEFECT LIST FORMAT** field specifies the format of the address descriptors in the defect list in the **FORMAT UNIT** parameter list.

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC 14776-323:2017

Table 34 defines support requirements for address descriptors based on the combinations of the FMTDATA bit, the Cmplst bit, the DEFECT LIST FORMAT field, and the DEFECT LIST LENGTH field.

Table 34 – FORMAT UNIT command address descriptor support requirements

Field in the FORMAT UNIT CDB			DEFECT LIST LENGTH ^a	Support	Comments
FMTDATA	Cmplst	DEFECT LIST FORMAT			
0	any	000b	Not available	M	Vendor specific defect information
1	0	Either: a) 000b (i.e., short block address descriptor)(see 6.2.2); b) 001b (i.e., extended bytes from index address descriptor)(see 6.2.3); c) 010b (i.e., extended physical sector address descriptor)(see 6.2.4); d) 011b (i.e., long block address descriptor)(see 6.2.5); e) 100b (i.e., bytes from index address descriptor)(see 6.2.6); or f) 101b (i.e., physical sector address descriptor)(see 6.2.7)	Zero	O	See ^b and ^d
	1			O	See ^b and ^e
	0		Non-zero	O	See ^c and ^d
	1			O	See ^c and ^e
	0	110b (i.e., vendor specific)	Zero	O	See ^b and ^d
			Non-zero	O	See ^c and ^d
			Zero	O	See ^b and ^e
			Non-zero	O	See ^c and ^e
All other combinations				Reserved	
^a This field is in the parameter list header. ^b No defect list is included in the parameter list. ^c A defect list is included in the parameter list. ^d The device server retains the existing GLIST. ^e The device server discards the existing GLIST.					

The CONTROL byte is defined in SAM-5.

If the LONGLIST bit is set to one in the FORMAT UNIT CDB, then the long parameter list header (see table 37) is used.

Table 37 – Long parameter list header

Byte	Bit	7	6	5	4	3	2	1	0
0		Reserved				PROTECTION FIELD USAGE			
		FOV	Format options bits				Obsolete	IMMED	Vendor specific
1			DPRY	DCRT	STPF	IP			
2		Reserved							
3		P_I_INFORMATION				PROTECTION INTERVAL EXPONENT			
4		(MSB)							
...		DEFECT LIST LENGTH							
7		(LSB)							

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC 14776-323:2017

The combination (see table 38) of the PROTECTION FIELD USAGE field and the FMTPINFO field (see 5.3.1) specifies the requested protection type (see 4.22.2).

Table 38 – FMTPINFO field and PROTECTION FIELD USAGE field (part 1 of 2)

Fields indicated by the device server		Fields specified by the application client		Description
SPT ^a	PROTECT ^b	FMTPINFO	PROTECTION FIELD USAGE	
xxx _b	0	00 _b	000 _b	The logical unit shall be formatted to type 0 protection ^c (see 4.22.2.2) resulting in the PROT_EN bit ^d being set to zero and the P_TYPE field ^d being set to 000 _b .
			>000 _b	Illegal ^e
		01 _b	xxx _b	Illegal ^f
		1x _b	xxx _b	Illegal ^f
xxx _b	1	00 _b	000 _b	The logical unit shall be formatted to type 0 protection ^c (see 4.22.2.2) resulting in the PROT_EN bit ^d being set to zero and the P_TYPE field ^d being set to 000 _b .
			>000 _b	Illegal ^e
xxx _b	1	01 _b	xxx _b	Illegal ^f
000 _b 001 _b 011 _b 111 _b	1	10 _b	000 _b	The logical unit shall be formatted to type 1 protection ^g (see 4.22.2.3) resulting in the PROT_EN bit ^d being set to one and the P_TYPE field ^d being set to 000 _b .
			>000 _b	Illegal ^e
010 _b 100 _b 101 _b	1	10 _b	xxx _b	Illegal ^f
000 _b	1	11 _b	xxx _b	Illegal ^f

^a See the Extended INQUIRY Data VPD page (see SPC-4) for the definition of the SPT field.

^b See the standard INQUIRY data (see SPC-4) for the definition of the PROTECT bit.

^c The device server shall format the medium to the logical block length specified in the mode parameter block descriptor of the mode parameter header (see SPC-4).

^d See the READ CAPACITY (16) parameter data (see 5.16.2) for the definition of the PROT_EN bit and P_TYPE field.

^e The device server shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST.

^f The device server shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

^g The device server shall format the medium to the logical block length specified in the mode parameter block descriptor of the mode parameter header plus eight bytes for each protection information interval (e.g., if the logical block length is 2 048 and the PROTECTION INTERVAL EXPONENT field set to two, then there are four 512 byte protection information intervals each followed by eight bytes of protection information resulting in a formatted logical block length of 2 080 bytes). Following a successful format operation, the PROT_EN bit in the READ CAPACITY (16) parameter data (see 5.16.2) indicates whether protection information (see 4.22) is enabled.

Table 38 – FMTPINFO field and PROTECTION FIELD USAGE field (part 2 of 2)

Fields indicated by the device server		Fields specified by the application client		Description
SPT ^a	PROTECT ^b	FMTPINFO	PROTECTION FIELD USAGE	
001b 010b 101b 111b	1	11b	000b	The logical unit shall be formatted to type 2 protection ^g (see 4.22.2.4) resulting in the PROT_EN bit ^d being set to one and the P_TYPE field ^d being set to 001b.
001b 010b	1	11b	>000b	Illegal ^e
011b 100b	1	11b	000b	Illegal ^e
011b 100b 101b 111b	1	11b	001b	The logical unit shall be formatted to type 3 protection ^g (see 4.22.2.5) resulting in the PROT_EN bit ^d being set to one and the P_TYPE field ^d being set to 010b.
			>001b	Illegal ^e
110b	1	10b 11b	xxx b	Reserved

^a See the Extended INQUIRY Data VPD page (see SPC-4) for the definition of the SPT field.

^b See the standard INQUIRY data (see SPC-4) for the definition of the PROTECT bit.

^c The device server shall format the medium to the logical block length specified in the mode parameter block descriptor of the mode parameter header (see SPC-4).

^d See the READ CAPACITY (16) parameter data (see 5.16.2) for the definition of the PROT_EN bit and P_TYPE field.

^e The device server shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST.

^f The device server shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

^g The device server shall format the medium to the logical block length specified in the mode parameter block descriptor of the mode parameter header plus eight bytes for each protection information interval (e.g., if the logical block length is 2 048 and the PROTECTION INTERVAL EXPONENT field set to two, then there are four 512 byte protection information intervals each followed by eight bytes of protection information resulting in a formatted logical block length of 2 080 bytes). Following a successful format operation, the PROT_EN bit in the READ CAPACITY (16) parameter data (see 5.16.2) indicates whether protection information (see 4.22) is enabled.

A format options valid (FOV) bit set to zero specifies that the device server shall use its default settings for the functionality represented by the DPROY bit, the DCRT bit, the STPF bit, and IP bit (i.e., the format options bits). If the FOV bit is set to zero, then the application client should set each of the format options bits to zero. If the FOV bit is set to one, and any of the format options bits are not set to zero, then the device server shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST.

A FOV bit set to one specifies that the device server shall process the format options bits as follows:

- a) a disable primary (DPROY) bit:
 - A) set to zero specifies that the device server shall not assign LBAs to parts of the medium identified as defective in the PLIST; or
 - B) set to one specifies that the device server shall not use the PLIST to identify defective areas of the medium, and the PLIST shall not be deleted;
- b) a disable certification (DCRT) bit:

- A) set to zero specifies that the device server shall perform a vendor specific medium certification operation and add address descriptors for defects that it detects during the certification operation to the GLIST; or
 - B) set to one specifies that the device server shall not perform any vendor specific medium certification process or format verification operation;
 - c) the stop format (STPF) bit controls the behavior of the device server if the device server has been requested to use the PLIST (i.e., the DPROY bit is set to zero) or the GLIST (i.e., the CMLST bit is set to zero) and one or more of the following occurs:
 - A) **list locate error**: the device server is not able to locate a specified defect list or determine whether a specified defect list exists; or
 - B) **list access error**: the device server encounters an error while accessing a specified defect list;
 - d) a STPF bit set to zero specifies that:
 - A) if a list locate error or a list access error occurs, then the device server shall continue to process the FORMAT UNIT command; and
 - B) after the format operation is complete, if:
 - a) a list locate error and a list access error both occurred, then the device server shall terminate the FORMAT UNIT command with CHECK CONDITION status at the completion of the command with the sense key set to RECOVERED ERROR and the additional sense code set to DEFECT LIST NOT FOUND or DEFECT LIST ERROR;
 - b) a list locate error occurred and a list access error did not occur, then the device server shall terminate the FORMAT UNIT command with CHECK CONDITION status with the sense key set to RECOVERED ERROR with the additional sense code set to DEFECT LIST NOT FOUND; and
 - c) a list access error occurred and a list locate error did not occur, then the device server shall terminate the FORMAT UNIT command with CHECK CONDITION status with the sense key set to RECOVERED ERROR with the additional sense code set to DEFECT LIST ERROR;
 - e) a STPF bit set to one specifies that:
 - A) if a list locate error occurs, then the device server shall terminate the FORMAT UNIT command with CHECK CONDITION status with the sense key set to MEDIUM ERROR and the additional sense code set to either DEFECT LIST NOT FOUND; or
 - B) if a list access error occurs, then the device server shall terminate the FORMAT UNIT command with CHECK CONDITION status with the sense key set to MEDIUM ERROR with the additional sense code set to DEFECT LIST ERROR;
- and
- f) an initialization pattern (IP) bit:
 - A) set to zero specifies that an initialization pattern descriptor (see 5.3.2.3) is not included and that the device server shall use its default initialization pattern; or
 - B) set to one specifies that:
 - a) an initialization pattern descriptor is included in the FORMAT UNIT parameter list following the parameter list header; and
 - b) if the device server does not support initialization pattern descriptors, then the device server shall terminate the FORMAT UNIT command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST.

An immediate (IMMED) bit set to zero specifies that the device server shall return status after the format operation has completed. An IMMED bit set to one specifies that the device server shall return status after the entire parameter list has been transferred.

The P__INFORMATION field, if any (i.e., if the long parameter list header is used), should be set to 0h. If the P__INFORMATION field is not set to zero, then the device server shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST.

For a type 0 or a type 1 protection information request, if the PROTECTION INTERVAL EXPONENT field, if any, is not set to 0h, then the device server shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST.

For a type 2 protection or a type 3 protection format request, the protection interval exponent determines the length of user data to be transferred before protection information is transferred (i.e., the protection information interval).

The protection information interval is calculated as follows:

$$\text{protection information interval} = \text{logical block length} \div 2^{(\text{protection interval exponent})}$$

where:

logical block length is the number of bytes of user data in a logical block (see 4.5)
 protection interval exponent is zero if the short parameter list header (see table 36) is used or the contents of the PROTECTION INTERVAL EXPONENT field if the long parameter list header (see table 37) is used

If the protection information interval calculates to a value that is not an even number (e.g., $520 \div 2^3 = 65$) or not a whole number (e.g., $520 \div 2^4 = 32.5$ and $520 \div 2^{10} = 0.508$), then the device server shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST.

The DEFECT LIST LENGTH field specifies the total length in bytes of the defect list (i.e., the address descriptors) that follow and does not include the length of the initialization pattern descriptor, if any. The formats for the address descriptor(s) are shown in 6.2.

5.3.2.3 Initialization pattern descriptor

The initialization pattern descriptor specifies that the device server initialize logical blocks to a specified pattern. The initialization pattern descriptor (see table 39) is transferred to the device server as part of the FORMAT UNIT parameter list.

Table 39 – Initialization pattern descriptor

Byte	Bit	7	6	5	4	3	2	1	0
0		Obsolete		SI	Reserved				
1		INITIALIZATION PATTERN TYPE							
2	(MSB)	INITIALIZATION PATTERN LENGTH (n - 3)							
3		(LSB)							
4		INITIALIZATION PATTERN							
...									
n									

A security initialize (SI) bit set to one specifies that the device server shall attempt to write the initialization pattern to all areas of the medium including those that may have been reassigned (i.e., are in a defect list). An SI bit set to zero specifies that the device server shall ignore:

- the FMTPINFO field;
- the FMTDATA bit;
- the CMLIST bit;
- the DEFECT LIST FORMAT field;
- all the bits and fields in the parameter list header, except the IMMED bit; and
- any defect list data.

The device server shall write the initialization pattern using a security erasure write technique. The security erasure write technique requirement and procedure is outside the scope of this standard. The device server is not required to write the initialization pattern over the header and other parts of the medium not previously accessible to the application client. If the device server is unable to write over any part of the medium that is

currently accessible to the application client or may be made accessible to the application client in the future (e.g., by clearing the defect list), then the device server shall terminate the command with CHECK CONDITION status with the sense key set to MEDIUM ERROR and the additional sense code set to the appropriate value for the condition. The device server shall attempt to rewrite all remaining parts of the medium even if some parts are not able to be rewritten.

NOTE 6 - The intent of the security erasure write is to render any previous user data unrecoverable by any analog or digital technique.

Migration from the SI bit to the SANITIZE command (see 5.24) is recommended for all implementations.

An SI bit set to zero specifies that the device server shall initialize the application client accessible part of the medium. The device server is not required to initialize other areas of the medium. The device server shall format the medium as defined in the FORMAT UNIT command.

The INITIALIZATION PATTERN TYPE field (see table 40) specifies the type of pattern the device server shall use to initialize each logical block within the application client accessible part of the medium. All bytes within a logical block shall be written with the initialization pattern.

Table 40 – INITIALIZATION PATTERN TYPE field

Code	Description
00h	Use a default initialization pattern ^a
01h	Repeat the pattern specified in the INITIALIZATION PATTERN field as required to fill the logical block ^b
02h to 7Fh	Reserved
80h to FFh	Vendor specific
^a If the INITIALIZATION PATTERN LENGTH field is not set to zero, then the device server shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST. ^b If the INITIALIZATION PATTERN LENGTH field is set to zero, then the device server shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST.	

The INITIALIZATION PATTERN LENGTH field specifies the number of bytes contained in the INITIALIZATION PATTERN field. If the initialization pattern length exceeds the current logical block length, then the device server shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST.

The INITIALIZATION PATTERN field specifies the initialization pattern.

5.4 GET LBA STATUS command

5.4.1 GET LBA STATUS command overview

The GET LBA STATUS command (see table 41) requests that the device server transfer parameter data describing the logical block provisioning status (see 4.7) for the specified LBA and zero or more subsequent LBAs to the Data-In Buffer.

The device server may or may not process this command as an uninterrupted sequence of actions (e.g., if concurrent operations are occurring that affect the logical block provisioning status, then the returned parameter data may be inconsistent or out of date).

This command uses the SERVICE ACTION IN (16) CDB format (see clause A.2).

Table 41 – GET LBA STATUS command

Byte	Bit	7	6	5	4	3	2	1	0
0		OPERATION CODE (9Eh)							
1		Reserved			SERVICE ACTION (12h)				
2	(MSB)	STARTING LOGICAL BLOCK ADDRESS							
...									
9									
10	(MSB)	ALLOCATION LENGTH							
...									
13									
14		Reserved							
15		CONTROL							

The OPERATION CODE field is defined in SPC-4 and shall be set to the value shown in table 41 for the GET LBA STATUS command.

The SERVICE ACTION field is defined in SPC-4 and shall be set to the value shown in table 41 for the GET LBA STATUS command.

The STARTING LOGICAL BLOCK ADDRESS field specifies the LBA of the first logical block addressed by this command. If the specified starting LBA exceeds the capacity of the medium (see 4.5), then the device server shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to LOGICAL BLOCK ADDRESS OUT OF RANGE.

The ALLOCATION LENGTH field is defined in SPC-4. In response to a GET LBA STATUS command, the device server may send less data to the Data-In Buffer than is specified by the allocation length. If, in response to a single GET LBA STATUS command, the device server does not send sufficient data to the Data-In Buffer to satisfy the requirement of the application client, then, to retrieve additional information, the application client may send additional GET LBA STATUS commands with different starting LBA values.

The CONTROL byte is defined in SAM-5.

5.4.2 GET LBA STATUS parameter data

5.4.2.1 GET LBA STATUS parameter data overview

The GET LBA STATUS parameter data (see table 42) contains an eight-byte header followed by one or more LBA status descriptors.

Table 42 – GET LBA STATUS parameter data

Byte	Bit	7	6	5	4	3	2	1	0
0	(MSB)	PARAMETER DATA LENGTH (n - 3)							(LSB)
...									
3									
4		Reserved							
...									
7									
LBA status descriptors									
8		LBA status descriptor [first] (see 5.4.2.2)							
...									
23									
		LBA status descriptor [last] (see 5.4.2.2) (if any)							
n - 15									
n									

The PARAMETER DATA LENGTH field indicates the number of bytes of parameter data that follow. The value in the PARAMETER DATA LENGTH field shall be:

- a) at least 20 (i.e., the available parameter data shall contain at least one LBA status descriptor); and
- b) four added to a multiple of 16 (i.e., the available parameter data shall end on a boundary between LBA Status descriptors).

Due to processing considerations outside the scope of this standard, two GET LBA STATUS commands with identical values in all CDB fields may result in two different values in the PARAMETER DATA LENGTH field.

The relationship between the PARAMETER DATA LENGTH field and the ALLOCATION LENGTH field in the CDB is defined in SPC-4.

5.4.2.2 LBA status descriptor

The LBA status descriptor (see table 43) contains LBA status information for one or more LBAs.

Table 43 – LBA status descriptor format

Byte	Bit	7	6	5	4	3	2	1	0
0	(MSB)	LBA STATUS LOGICAL BLOCK ADDRESS							
...									
7									
8	(MSB)	NUMBER OF LOGICAL BLOCKS							
...									
11									
12		Reserved			PROVISIONING STATUS				
13		Reserved							
...									
15									

The LBA STATUS LOGICAL BLOCK ADDRESS field contains the first LBA of the LBA extent for which this descriptor reports LBA status.

The NUMBER OF LOGICAL BLOCKS field contains the number of logical blocks in that LBA extent. The device server should return the largest possible value in the NUMBER OF LOGICAL BLOCKS field.

The PROVISIONING STATUS field is defined in table 44.

Table 44 – PROVISIONING STATUS field

Code	Description
0h	Each LBA in the LBA extent is mapped (see 4.7.4.6) or has an unknown state.
1h	Each LBA in the LBA extent is deallocated (see 4.7.4.7).
2h	Each LBA in the LBA extent is anchored (see 4.7.4.8).
All others	Reserved

If the logical unit is fully provisioned (see 4.7.2), then the PROVISIONING STATUS field for all LBAs shall be set to 0h (i.e., mapped or unknown).

5.4.2.3 LBA status descriptor relationships

The LBA STATUS LOGICAL BLOCK ADDRESS field in the first LBA status descriptor returned in the GET LBA STATUS parameter data shall contain the value specified in the STARTING LOGICAL BLOCK ADDRESS field of the CDB. For subsequent LBA status descriptors, the contents of the LBA STATUS LOGICAL BLOCK ADDRESS field shall contain the sum of the values in:

- the LBA STATUS LOGICAL BLOCK ADDRESS field in the previous LBA status descriptor; and
- the NUMBER OF LOGICAL BLOCKS field in the previous LBA status descriptor.

Adjacent LBA status descriptors may have the same values for the PROVISIONING STATUS field.

5.5 ORWRITE (16) command

The ORWRITE (16) command (see table 45) requests that the device server perform an ORWRITE command set operation (see 4.29.4).

Table 45 – ORWRITE (16) command

Byte	Bit	7	6	5	4	3	2	1	0
0		OPERATION CODE (8Bh)							
1		ORPROTECT			DPO	FUA	Reserved		
2	(MSB)	LOGICAL BLOCK ADDRESS							
...									
9		(LSB)							
10	(MSB)	TRANSFER LENGTH							
...									
13		(LSB)							
14		Reserved			GROUP NUMBER				
15		CONTROL							

The OPERATION CODE field is defined in SPC-4 and shall be set to the value shown in table 45 for the ORWRITE (16) command.

See the READ (10) command (see 5.11) for the definition of the FUA bit specifying behavior for read operations. See the WRITE (10) command (see 5.33) for the definition of the FUA bit specifying behavior for write operations. See the READ (10) command (see 5.11) for the definition of the DPO bit. See the PRE-FETCH (10) command (see 5.8) for the definition of the LOGICAL BLOCK ADDRESS field. See the PRE-FETCH (10) command (see 5.8) and 4.23 for the definition of the GROUP NUMBER field.

The TRANSFER LENGTH field specifies the number of contiguous logical blocks of data that are read, transferred from the Data-Out Buffer, and ORed into a bitmap buffer, starting with the logical block referenced by the LBA specified by the LOGICAL BLOCK ADDRESS field. If the specified LBA and the specified transfer length exceed the capacity of the medium (see 4.5), then the device server shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to LOGICAL BLOCK ADDRESS OUT OF RANGE. The TRANSFER LENGTH field is constrained by the MAXIMUM TRANSFER LENGTH field in the Block Limits VPD page (see 6.6.3).

The CONTROL byte is defined in SAM-5.

The device server shall:

- a) check protection information from the read operations based on the ORPROTECT field as described in table 46; and
- b) check protection information transferred from the Data-Out Buffer based on the ORPROTECT field as described in table 47.

The order of the user data and protection information checks and comparisons is vendor specific.

The device server shall check the protection information from the read operations based on the ORPROTECT field as described in table 46. All footnotes for table 46 are at the end of the table.

Table 46 – ORPROTECT field - checking protection information from the read operations (part 1 of 3)

Code	Logical unit formatted with protection information	Field in protection information ^g	Extended INQUIRY Data VPD page bit value ^f	If check fails ^{d e} , additional sense code
000b	Yes ^{i j}	LOGICAL BLOCK GUARD	GRD_CHK = 1	LOGICAL BLOCK GUARD CHECK FAILED
			GRD_CHK = 0	No check performed
		LOGICAL BLOCK APPLICATION TAG	APP_CHK = 1 ^c	LOGICAL BLOCK APPLICATION TAG CHECK FAILED
			APP_CHK = 0	No check performed
		LOGICAL BLOCK REFERENCE TAG	REF_CHK = 1 ^h	LOGICAL BLOCK REFERENCE TAG CHECK FAILED
			REF_CHK = 0	No check performed
No	No protection information on the medium to check.			
001b 101b ^b	Yes	LOGICAL BLOCK GUARD	GRD_CHK = 1	LOGICAL BLOCK GUARD CHECK FAILED
			GRD_CHK = 0	No check performed
		LOGICAL BLOCK APPLICATION TAG	APP_CHK = 1 ^c	LOGICAL BLOCK APPLICATION TAG CHECK FAILED
			APP_CHK = 0	No check performed
		LOGICAL BLOCK REFERENCE TAG	REF_CHK = 1 ^h	LOGICAL BLOCK REFERENCE TAG CHECK FAILED
			REF_CHK = 0	No check performed
No	Error condition ^a			
010b ^b	Yes	LOGICAL BLOCK GUARD	No check performed	
		LOGICAL BLOCK APPLICATION TAG	APP_CHK = 1 ^c	LOGICAL BLOCK APPLICATION TAG CHECK FAILED
			APP_CHK = 0	No check performed
		LOGICAL BLOCK REFERENCE TAG	REF_CHK = 1 ^h	LOGICAL BLOCK REFERENCE TAG CHECK FAILED
			REF_CHK = 0	No check performed
		No	Error condition ^a	

Table 46 – ORPROTECT field - checking protection information from the read operations (part 2 of 3)

Code	Logical unit formatted with protection information	Field in protection information ^g	Extended INQUIRY Data VPD page bit value ^f	If check fails ^{d e} , additional sense code
011b ^b	Yes	LOGICAL BLOCK GUARD	No check performed	
		LOGICAL BLOCK APPLICATION TAG	No check performed	
		LOGICAL BLOCK REFERENCE TAG	No check performed	
	No	Error condition ^a		
100b ^b	Yes	LOGICAL BLOCK GUARD	GRD_CHK = 1	LOGICAL BLOCK GUARD CHECK FAILED
			GRD_CHK = 0	No check performed
		LOGICAL BLOCK APPLICATION TAG	No check performed	
		LOGICAL BLOCK REFERENCE TAG	No check performed	
	No	Error condition ^a		

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC 14776-323:2017

Table 46 – ORPROTECT field - checking protection information from the read operations (part 3 of 3)

Code	Logical unit formatted with protection information	Field in protection information ^g	Extended INQUIRY Data VPD page bit value ^f	If check fails ^{d e} , additional sense code
110b to 111b	Reserved			
<p>^a If the logical unit supports protection information (see 4.22) and has not been formatted with protection information, then the device server shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.</p> <p>^b If the logical unit does not support protection information, then the device server should terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.</p> <p>^c If the device server has knowledge of the contents of the logical block application tag field, then the device server shall check the logical block application tag. If the ato bit in the Control mode page (see SPC-4) is set to one, then this knowledge is acquired from:</p> <p>a) the Application Tag mode page (see 6.5.3), if the atmpe bit in the Control mode page (see SPC-4) is set to one; or</p> <p>b) a method not defined by this standard, if the ATMPE bit is set to zero.</p> <p>^d If the device server terminates the command with CHECK CONDITION status, then the device server shall set the sense key to ABORTED COMMAND.</p> <p>^e If multiple errors occur while the device server is processing the command, then the selection by the device server of which error to report is not defined by this standard.</p> <p>^f See the Extended INQUIRY Data VPD page (see SPC-4) for the definitions of the GRD_CHK bit, the APP_CHK bit, and the REF_CHK bit.</p> <p>^g If the device server detects:</p> <p>a) a LOGICAL BLOCK APPLICATION TAG field set to FFFFh, and type 1 protection (see 4.22.2.3) is enabled; or</p> <p>b) a LOGICAL BLOCK APPLICATION TAG field set to FFFFh, the LOGICAL BLOCK REFERENCE TAG field set to FFFF_FFFFh, and type 3 protection (see 4.22.2.5) is enabled, then the device server shall not check any protection information in the associated protection information interval.</p> <p>^h If type 1 protection is enabled, then the device server shall check the logical block reference tag by comparing it to the lower four bytes of the LBA associated with the logical block. If type 3 protection is enabled, then the device server shall check each logical block reference tag only if the device server has knowledge of the contents of the LOGICAL BLOCK REFERENCE TAG field. The method for acquiring this knowledge is not defined by this standard.</p> <p>ⁱ If the RWWP bit in the Control mode page (see SPC-4) is set to one, then the device server shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.</p> <p>^j If the DPICZ bit in the Control mode page (see SPC-4) is set to one, and the RWWP bit in the Control mode page is set to zero, then protection information shall not be checked.</p>				

The device server shall check the protection information transferred from the Data-Out Buffer based on the ORPROTECT field as described in table 47. All footnotes for table 47 are at the end of the table.

Table 47 – ORPROTECT field - checking protection information from the Data-Out Buffer (part 1 of 2)

Code	Logical unit formatted with protection information	Field in protection information	Device server check	If check fails ^{d h} , additional sense code
000b	Yes ^{e f g}	No protection information in the Data-Out buffer to check		
	No	No protection information in the Data-Out buffer to check		
001b ^b	Yes	LOGICAL BLOCK GUARD	Shall	LOGICAL BLOCK GUARD CHECK FAILED
		LOGICAL BLOCK APPLICATION TAG	Dependent on RWWP ^c	LOGICAL BLOCK APPLICATION TAG CHECK FAILED
		LOGICAL BLOCK REFERENCE TAG	Shall (except for type 3) ⁱ	LOGICAL BLOCK REFERENCE TAG CHECK FAILED
	No	Error condition ^a		
010b ^b	Yes	LOGICAL BLOCK GUARD	Shall not	No check performed
		LOGICAL BLOCK APPLICATION TAG	Dependent on RWWP ^c	LOGICAL BLOCK APPLICATION TAG CHECK FAILED
		LOGICAL BLOCK REFERENCE TAG	May ⁱ	LOGICAL BLOCK REFERENCE TAG CHECK FAILED
	No	Error condition ^a		
011b ^b	Yes	LOGICAL BLOCK GUARD	Shall not	No check performed
		LOGICAL BLOCK APPLICATION TAG	Shall not	No check performed
		LOGICAL BLOCK REFERENCE TAG	Shall not	No check performed
	No	Error condition ^a		
100b ^b	Yes	LOGICAL BLOCK GUARD	Shall	LOGICAL BLOCK GUARD CHECK FAILED
		LOGICAL BLOCK APPLICATION TAG	Shall not	No check performed
		LOGICAL BLOCK REFERENCE TAG	Shall not	No check performed
	No	Error condition ^a		

Table 47 – ORPROTECT field - checking protection information from the Data-Out Buffer (part 2 of 2)

Code	Logical unit formatted with protection information	Field in protection information	Device server check	If check fails ^{d h} , additional sense code
101b ^b	Yes	LOGICAL BLOCK GUARD	Shall	LOGICAL BLOCK GUARD CHECK FAILED
		LOGICAL BLOCK APPLICATION TAG	Dependent on RWWP ^c	LOGICAL BLOCK APPLICATION TAG CHECK FAILED
		LOGICAL BLOCK REFERENCE TAG	May ⁱ	LOGICAL BLOCK REFERENCE TAG CHECK FAILED
	No	Error condition ^a		
110b to 111b	Reserved			

^a If a logical unit supports protection information (see 4.22) and has not been formatted with protection information, then the device server shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

^b If the logical unit does not support protection information, then the device server should terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

^c If the ATO bit is set to one in the Control mode page (see SPC-4), and the device server has knowledge of the contents of the LOGICAL BLOCK APPLICATION TAG field, then the device server:

a) may check each logical block application tag if the RWWP bit is set to zero in the Control mode page (see SPC-4); and

b) shall check each logical block application tag if the RWWP bit is set to one in the Control mode page. If the ATMPE bit in the Control mode page (see SPC-4) is set to one, then this knowledge is acquired from the Application Tag mode page. If the ATMPE bit is set to zero, then the method for acquiring this knowledge is not defined by this standard.

^d If the device server terminates the command with CHECK CONDITION status, then the device server shall set the sense key to ABORTED COMMAND.

^e The device server shall write a generated CRC (see 4.22.4.2) into each LOGICAL BLOCK GUARD field.

^f If the RWWP bit in the Control mode page (see SPC-4) is set to one, then the device server shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB. If the RWWP bit is set to zero, and:

a) type 1 protection is enabled, then the device server shall write the least significant four bytes of each LBA into the LOGICAL BLOCK REFERENCE TAG field of each of the written logical blocks; or

b) type 3 protection is enabled, then the device server shall write a value of FFFF_FFFFh into the LOGICAL BLOCK REFERENCE TAG field of each of the written logical blocks.

^g If the ATO bit is set to one in the Control mode page (see SPC-4), then the device server shall write FFFFh into each LOGICAL BLOCK APPLICATION TAG field. If the ATO bit is set to zero, then the device server may write any value into each LOGICAL BLOCK APPLICATION TAG field.

^h If multiple errors occur while the device server is processing the command, then the selection by the device server of which error to report is not defined by this standard.

ⁱ If type 1 protection is enabled, then the device server shall check the logical block reference tag by comparing it to the lower four bytes of the LBA associated with the logical block. If type 3 protection is enabled, then the device server may check each logical block reference tag if the ATO bit is set to one in the Control mode page (see SPC-4), and the device server has knowledge of the contents of the LOGICAL BLOCK REFERENCE TAG. The method for acquiring this knowledge is not defined by this standard.

5.6 ORWRITE (32) command

The ORWRITE (32) command (see table 48) requests that the device server perform one of the following ORWRITE command (see 4.29) operations:

- a) a change generation and clear operation (see 4.29.3); or
- b) a set operation (see 4.29.4).

Table 48 – ORWRITE (32) command

Byte	Bit	7	6	5	4	3	2	1	0	
0	OPERATION CODE (7Fh)									
1	CONTROL									
2	Reserved						BMOP			
3	Reserved				PREVIOUS GENERATION PROCESSING					
4	Reserved									
5	Reserved									
6	Reserved				GROUP NUMBER					
7	ADDITIONAL CDB LENGTH (18h)									
8	(MSB)	SERVICE ACTION (000Eh)							(LSB)	
9										
10	ORPROTECT				DPO	FUA	Reserved			
11	Reserved									
12	(MSB)	LOGICAL BLOCK ADDRESS							(LSB)	
19										
20	(MSB)	EXPECTED ORWGENERATION							(LSB)	
23										
24	(MSB)	NEW ORWGENERATION							(LSB)	
27										
28	(MSB)	TRANSFER LENGTH							(LSB)	
31										

The OPERATION CODE field, the ADDITIONAL CDB LENGTH field, and the SERVICE ACTION field are defined in SPC-4 and shall be set to the values shown in table 48 for the ORWRITE (32) command.

The CONTROL byte is defined in SAM-5.

The bitmap operation (BMOP) field specifies the operation as described in table 49.

Table 49 – BMOP field

Code	Description
000b	The device server shall perform a set operation (see 4.29.4), and the contents of the PREVIOUS GENERATION PROCESSING field and NEW ORWGENERATION field shall be ignored.
001b	The device server shall perform a change generation and clear operation (see 4.29.3).
All others	Reserved

The PREVIOUS GENERATION PROCESSING field specifies the policy for performing future set operations that is to be established in the device server by a successful change generation and clear operation (see 4.29.2.2).

See the ORWRITE (16) command (see 5.5) for the definitions of the FUA bit, the DPO bit, the ORPROTECT field, the LOGICAL BLOCK ADDRESS field, the TRANSFER LENGTH field, and the GROUP NUMBER field.

The EXPECTED ORWGENERATION field contains a code that is compared with generation codes established and maintained by the device server.

The NEW ORWGENERATION field specifies the current ORWgeneration code that is to be established in the device server by a successful change generation and clear operation (see 4.29.3).

The device server shall:

- a) check protection information from the read operations based on the ORPROTECT field as described in table 46; and
- b) check protection information transferred from the Data-Out Buffer based on the ORPROTECT field as described in table 47.

The order of the user data and protection information checks and comparisons is vendor specific.

5.7 POPULATE TOKEN command

5.7.1 POPULATE TOKEN command overview

The POPULATE TOKEN command (see table 50) requests that the copy manager (see SPC-4) create a point in time ROD token that represents the specified logical blocks (see 4.30).

Each logical block represented by the point in time ROD token includes logical block data.

Table 50 – POPULATE TOKEN command

Byte	Bit	7	6	5	4	3	2	1	0
0		OPERATION CODE (83h)							
1		Reserved			SERVICE ACTION (10h)				
2		Reserved							
...									
5									
6	(MSB)	LIST IDENTIFIER							
...									
9									
10	(MSB)	PARAMETER LIST LENGTH							
...									
13									
14		Reserved			GROUP NUMBER				
15		CONTROL							

The OPERATION CODE field and the SERVICE ACTION field are defined in SPC-4 and shall be set to the values shown in table 50 for the POPULATE TOKEN command.

The LIST IDENTIFIER field is defined in SPC-4. The list identifier shall be processed as if the LIST ID USAGE field in the parameter data for an EXTENDED COPY(LID4) command (see SPC-4) is set to 00b.

The PARAMETER LIST LENGTH field specifies the length in bytes of the parameter list that is available to be transferred from the Data-Out Buffer. If the parameter list length is greater than zero and less than 00000010h (i.e., 16), then the copy manager shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to PARAMETER LIST LENGTH ERROR. A PARAMETER LIST LENGTH field set to zero specifies that no data shall be transferred. This shall not be considered an error.

See the PRE-FETCH (10) command (see 5.8) and 4.23 for the definition of the GROUP NUMBER field.

The CONTROL byte is defined in SAM-5.

5.7.2 POPULATE TOKEN parameter list

The parameter list for the POPULATE TOKEN command is shown in table 51.

Table 51 – POPULATE TOKEN parameter list

Byte	Bit	7	6	5	4	3	2	1	0	
0	(MSB)	POPULATE TOKEN DATA LENGTH (n - 1)								
1									(LSB)	
2		Reserved						RTV	IMMED	
3		Reserved								
4	(MSB)	INACTIVITY TIMEOUT								
...										
7									(LSB)	
8	(MSB)	ROD TYPE								
...										
11									(LSB)	
12		Reserved								
13										
14	(MSB)	BLOCK DEVICE RANGE DESCRIPTOR LENGTH (n - 15)								
15									(LSB)	
Block device range descriptor list										
16		Block device range descriptor [first] (see 5.7.3)								
...										
31										
		⋮								
n - 15		Block device range descriptor [last] (see 5.7.3) (if any)								
...										
n										

The POPULATE TOKEN DATA LENGTH field specifies the length in bytes of the data that is available to be transferred from the Data-Out Buffer. The populate token data length does not include the number of bytes in the POPULATE TOKEN DATA LENGTH field. If the POPULATE TOKEN DATA LENGTH field is less than 001Eh (i.e., 30), then the copy manager shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST.

A ROD type valid (RTV) bit set to zero specifies that the copy manager may create a ROD token with any point in time copy ROD type and shall ignore the contents of the ROD TYPE field. An RTV bit set to one specifies that the copy manager shall use the contents of the ROD TYPE field to create the point in time copy ROD.

The immediate (IMMED) bit specifies when the copy manager shall return status for the POPULATE TOKEN command. If the IMMED bit is set to zero, then the copy manager shall process the POPULATE TOKEN command until all specified operations are complete or an error is detected. If the IMMED bit is set to one, then the copy manager:

- 1) shall validate the CDB (i.e., detect and report all errors in the CDB);
- 2) shall transfer all the parameter data to the copy manager;
- 3) may validate the parameter data;

- 4) shall complete the POPULATE TOKEN command with GOOD status; and
- 5) shall complete performing of all specified operations as a background operation (see SPC-4).

If the INACTIVITY TIMEOUT field is not set to zero, then the INACTIVITY TIMEOUT field specifies the number of seconds that the copy manager should wait for the next third-party copy command that uses the ROD token created by the processing of this command before the copy manager invalidates that ROD token due to expiration of this timeout (see SPC-4). If the INACTIVITY TIMEOUT field is set to a value larger than the value in the MAXIMUM INACTIVITY TIMEOUT field in the Block Device ROD Token Limits descriptor (see 6.6.6.3), then the copy manager shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST.

If the INACTIVITY TIMEOUT field is set to zero, then the DEFAULT INACTIVITY TIMEOUT field in the Block Device ROD Token Limits descriptor (see 6.6.6.3) specifies the number of seconds that the copy manager should wait for the next third-party copy command (see 6.6.6) that uses the ROD token created by the processing of this command before the copy manager invalidates that ROD token due to expiration of this timeout.

If the RTV bit is set to one, then the ROD TYPE field specifies the ROD type (see SPC-4) for creating the point in time copy ROD token. The copy manager shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST, if:

- a) the copy manager does not support the specified ROD type for use with the POPULATE TOKEN command; or
- b) the ROD TYPE field specifies a ROD type (see SPC-4) that is not a point in time copy ROD.

The BLOCK DEVICE RANGE DESCRIPTOR LENGTH field specifies the length in bytes of the block device range descriptor list. The block device range descriptor list length should be a multiple of 16. If the block device range descriptor list length is not a multiple of 16, then the last block device range descriptor is incomplete and shall be ignored. If the block device range descriptor list length is less than 0010h (i.e., 16), then the copy manager shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST.

If any block device range descriptors in the block device range descriptor list are truncated due to the parameter list length in the CDB, then those block device range descriptors shall be ignored.

If the number of complete block device range descriptors is larger than the maximum range descriptors value in the Block Device ROD Token Limits descriptor (see 6.6.6.3), then the copy manager shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to TOO MANY SEGMENT DESCRIPTORS.

If the same LBA is included in more than one block device range descriptor, then the copy manager shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST.

If the number of bytes of user data represented by the sum of the contents of the NUMBER OF LOGICAL BLOCKS fields in all of the complete block device range descriptors is larger than:

- a) the MAXIMUM BYTES IN BLOCK ROD field in the block ROD device type specific features descriptor in the ROD token features third-party copy descriptor in the Third-party Copy VPD page (see SPC-4) and that field is set to a nonzero value; or
- b) the MAXIMUM TOKEN TRANSFER SIZE field in the Block Device ROD Token Limits descriptor (see 6.6.6.3) and that field is set to a nonzero value,

then the copy manager shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST.

5.7.3 Block device range descriptor

The block device range descriptor is defined in table 52.

Table 52 – Block device range descriptor

Byte	Bit	7	6	5	4	3	2	1	0
0	(MSB)	LOGICAL BLOCK ADDRESS							
...									
7									
8	(MSB)	NUMBER OF LOGICAL BLOCKS							
...									
11									
12		Reserved							
...									
15									

The LOGICAL BLOCK ADDRESS field specifies the first LBA on which the copy manager shall operate for this block device range descriptor.

The NUMBER OF LOGICAL BLOCKS field specifies the number of logical blocks on which the copy manager shall operate for this block device range descriptor beginning with the LBA specified by the LOGICAL BLOCK ADDRESS field.

Processing of block device range descriptors with a number of logical blocks that is not a multiple of the OPTIMAL BLOCK ROD LENGTH GRANULARITY field in the block ROD device type specific features descriptor in the ROD token features third-party copy descriptor in the Third-party Copy VPD page (see SPC-4) may incur delays in processing. If the OPTIMAL BLOCK ROD LENGTH GRANULARITY field in the block ROD device type specific features descriptor in the ROD token features third-party copy descriptor in the Third-party Copy VPD page is not reported, then the optimal transfer length granularity in the Block Limits VPD page (see 6.6.3) may indicate the granularity.

For a POPULATE TOKEN command, processing of block device range descriptors where the number of bytes of user data contained in the number of logical blocks exceeds the OPTIMAL BYTES TO TOKEN PER SEGMENT field in the block ROD device type specific features descriptor in the ROD token features third-party copy descriptor in the Third-party Copy VPD page may incur delays in processing.

For a WRITE USING TOKEN command (see 5.46), processing of block device range descriptors where the number of bytes of user data contained in the number of logical blocks exceeds the OPTIMAL BYTES FROM TOKEN PER SEGMENT field in the block ROD device type specific features descriptor in the ROD token features third-party copy descriptor in the Third-party Copy VPD page may incur delays in processing.

If the number of bytes of user data contained in the number of logical blocks is greater than:

- 1) the value in the MAXIMUM BYTES IN BLOCK ROD field in the block ROD device type specific features descriptor in the ROD token features third-party copy descriptor in the Third-party Copy VPD page, and the MAXIMUM BYTES IN BLOCK ROD field is set to a nonzero value; or
- 2) the value in the MAXIMUM TRANSFER LENGTH field in the Block Limits VPD page (see 6.6.3), the MAXIMUM TRANSFER LENGTH field is set to a nonzero value, and the MAXIMUM BYTES IN BLOCK ROD field in the block ROD device type specific features descriptor is not reported,

then the copy manager shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST.

If the NUMBER OF LOGICAL BLOCKS field is set to zero, then the copy manager shall perform no operation for this block device range descriptor. This condition shall not be considered an error.

If the specified LBA and the specified number of logical blocks exceed the capacity of the medium (see 4.5), then the copy manager shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to LOGICAL BLOCK ADDRESS OUT OF RANGE.

5.8 PRE-FETCH (10) command

The PRE-FETCH (10) command (see table 53) requests that the device server:

- a) for any mapped LBAs specified by the command that are not already contained in cache, perform read medium operations and write cache operations (see 4.15); and
- b) for any unmapped LBAs specified by the command, update the volatile cache and/or non-volatile cache to prevent retrieval of stale data.

No data shall be transferred to the Data-In Buffer.

Migration from the PRE-FETCH (10) command to the PRE-FETCH (16) command is recommended for all implementations.

Table 53 – PRE-FETCH (10) command

Byte	Bit	7	6	5	4	3	2	1	0	
0		OPERATION CODE (34h)								
1		Reserved						IMMED	Obsolete	
2	(MSB)	LOGICAL BLOCK ADDRESS								
...										
5		(LSB)								
6		Reserved			GROUP NUMBER					
7	(MSB)	PREFETCH LENGTH								
8										(LSB)
9		CONTROL								

The OPERATION CODE field is defined in SPC-4 and shall be set to the value shown in table 53 for the PRE-FETCH (10) command.

An immediate (IMMED) bit set to zero specifies that status shall be returned after the operation is complete. An IMMED bit set to one specifies that the device server shall:

- a) validate the CDB;
- b) if the cache has:
 - A) sufficient capacity to accept all of the specified logical blocks, then complete the command with CONDITION MET status; or
 - B) insufficient capacity to accept all of the specified logical blocks, then complete the command with GOOD status;
 and
- c) if one or more of the specified logical blocks are not successfully transferred to the cache for reasons other than lack of cache capacity, then report a deferred error (see SPC-4).

The LOGICAL BLOCK ADDRESS field specifies the LBA of the first logical block (see 4.5) accessed by this command.

A GROUP NUMBER field set to a non-zero value specifies the group into which attributes associated with the command should be collected (see 4.23). A GROUP NUMBER field set to zero specifies that any attributes associated with the command shall not be collected into any group.

The PREFETCH LENGTH field specifies the number of contiguous logical blocks that shall be pre-fetched (i.e., transferred to the cache from the medium), starting with the LBA specified by the LOGICAL BLOCK ADDRESS field. A PREFETCH LENGTH field set to zero specifies that all logical blocks starting with the LBA specified in the LOGICAL BLOCK ADDRESS field to the last logical block on the medium shall be pre-fetched. Any other value specifies the number of logical blocks that shall be pre-fetched. If the specified LBA and the specified prefetch length exceed the capacity of the medium (see 4.5), then the device server shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to LOGICAL BLOCK ADDRESS OUT OF RANGE.

The CONTROL byte is defined in SAM-5.

If the IMMED bit is set to zero, and the specified logical blocks were transferred to the cache without error, then the device server shall complete the command with CONDITION MET status.

If the IMMED bit is set to zero and the cache does not have sufficient capacity to accept all of the specified logical blocks, then the device server shall transfer to the cache as many of the specified logical blocks that fit. If these logical blocks are transferred without error, then the device server shall complete the command with GOOD status.

5.9 PRE-FETCH (16) command

The PRE-FETCH (16) command (see table 54) requests that the device server perform the actions defined for the PRE-FETCH (10) command (see 5.8).

No data shall be transferred to the Data-In Buffer.

Table 54 – PRE-FETCH (16) command

Byte	Bit	7	6	5	4	3	2	1	0	
0		OPERATION CODE (90h)								
1		Reserved						IMMED	Reserved	
2	(MSB)	LOGICAL BLOCK ADDRESS								
...										
9										(LSB)
10	(MSB)	PREFETCH LENGTH								
...										
13										(LSB)
14		Reserved				GROUP NUMBER				
15		CONTROL								

The OPERATION CODE field is defined in SPC-4 and shall be set to the value shown in table 54 for the PRE-FETCH (16) command.

See the PRE-FETCH (10) command (see 5.8) for the definitions of the other fields in this command.

5.10 PREVENT ALLOW MEDIUM REMOVAL command

The PREVENT ALLOW MEDIUM REMOVAL command (see table 55) requests that the logical unit enable or disable the removal of the medium. If medium removal is prevented on any I_T nexus that has access to the logical unit, then the logical unit shall not allow medium removal.

Table 55 – PREVENT ALLOW MEDIUM REMOVAL command

Byte	Bit	7	6	5	4	3	2	1	0
0	OPERATION CODE (1Eh)								
1	Reserved								
2	Reserved								
3	Reserved								
4	Reserved							PREVENT	
5	CONTROL								

The OPERATION CODE field is defined in SPC-4 and shall be set to the value shown in table 55 for the PREVENT ALLOW MEDIUM REMOVAL command.

Table 56 defines the PREVENT field.

Table 56 – PREVENT field

Value	Description
00b	Medium removal is allowed.
01b	Medium removal shall be prevented.
10b to 11b	Obsolete

The CONTROL byte is defined in SAM-5.

The prevention of medium removal shall begin when any application client issues a PREVENT ALLOW MEDIUM REMOVAL command with the PREVENT field set to 01b (i.e., medium removal prevented). The prevention of medium removal for the logical unit shall no longer be prevented after:

- a) one of the following occurs for each I_T nexus through which medium removal had been prevented:
 - A) receipt of a PREVENT ALLOW MEDIUM REMOVAL command with the PREVENT field set to 00b; or
 - B) an I_T nexus loss;
- b) a power on;
- c) a hard reset; or
- d) a logical unit reset.

If possible, the device server shall perform a synchronize cache operation before ending the prevention of medium removal.

If a persistent reservation or registration is being preempted by a PERSISTENT RESERVE OUT command with PREEMPT AND ABORT service action (see SPC-4) or PERSISTENT RESERVE OUT command with CLEAR service action (see SPC-4), then the equivalent of a PREVENT ALLOW MEDIUM REMOVAL command with the PREVENT field set to 00b shall be processed for each the I_T nexuses associated with the persistent reservation or registrations being preempted allowing an application client to override the prevention of medium removal function for a SCSI initiator port (e.g., an initiator port is not operating correctly).

While a prevention of medium removal condition is in effect, the logical unit shall inhibit mechanisms that allow removal of the medium by an operator.

5.11 READ (10) command

The READ (10) command (see table 57) requests that the device server:

- a) perform read operations from the specified LBAs: and
- b) transfer the requested logical block data to the Data-In Buffer.

The logical block data transferred to the Data-In Buffer shall include protection information based on the value in the RDPROTECT field (see table 58) and the medium format.

Migration from the READ (10) command to the READ (16) command is recommended for all implementations.

Table 57 – READ (10) command

Byte	Bit	7	6	5	4	3	2	1	0
0		OPERATION CODE (28h)							
1		RDPROTECT			DPO	FUA	RARC	Obsolete	Obsolete
2	(MSB)	LOGICAL BLOCK ADDRESS							
...									
5		(LSB)							
6		Reserved			GROUP NUMBER				
7	(MSB)	TRANSFER LENGTH							
8									
9		CONTROL							

The OPERATION CODE field is defined in SPC-4 and shall be set to the value shown in table 57 for the READ (10) command.

The device server shall check the protection information from the read operations before returning status for the command based on the RDPROTECT field as described in table 58. All footnotes for table 58 are at the end of the table.

Table 58 – RDPROTECT field (part 1 of 3)

Code	Logical unit formatted with protection information	Shall device server transmit protection information?	Field in protection information ^h	Extended INQUIRY Data VPD page bit value ^g	If check fails ^{d,f} , additional sense code
000b	Yes ^j	No	LOGICAL BLOCK GUARD	GRD_CHK = 1	LOGICAL BLOCK GUARD CHECK FAILED
				GRD_CHK = 0	No check performed
			LOGICAL BLOCK APPLICATION TAG	APP_CHK = 1 ^c	LOGICAL BLOCK APPLICATION TAG CHECK FAILED
				APP_CHK = 0	No check performed
			LOGICAL BLOCK REFERENCE TAG	REF_CHK = 1 ⁱ	LOGICAL BLOCK REFERENCE TAG CHECK FAILED
				REF_CHK = 0	No check performed
No	No protection information available to check				
001b 101b ^b	Yes	Yes ^e	LOGICAL BLOCK GUARD	GRD_CHK = 1	LOGICAL BLOCK GUARD CHECK FAILED
				GRD_CHK = 0	No check performed
			LOGICAL BLOCK APPLICATION TAG	APP_CHK = 1 ^c	LOGICAL BLOCK APPLICATION TAG CHECK FAILED
				APP_CHK = 0	No check performed
			LOGICAL BLOCK REFERENCE TAG	REF_CHK = 1 ⁱ	LOGICAL BLOCK REFERENCE TAG CHECK FAILED
				REF_CHK = 0	No check performed
No ^a	No protection information available to transmit to the Data-In Buffer or for checking				

Table 58 – RDPROTECT field (part 2 of 3)

Code	Logical unit formatted with protection information	Shall device server transmit protection information?	Field in protection information ^h	Extended INQUIRY Data VPD page bit value ^g	If check fails ^{d,f} , additional sense code
010b ^b	Yes	Yes ^e	LOGICAL BLOCK GUARD	No check performed	
			LOGICAL BLOCK APPLICATION TAG	APP_CHK = 1 ^c	LOGICAL BLOCK APPLICATION TAG CHECK FAILED
				APP_CHK = 0	No check performed
			LOGICAL BLOCK REFERENCE TAG	REF_CHK = 1 ⁱ	LOGICAL BLOCK REFERENCE TAG CHECK FAILED
	REF_CHK = 0	No check performed			
No ^a	No protection information available to transmit to the Data-In Buffer or for checking				
011b ^b	Yes	Yes ^e	LOGICAL BLOCK GUARD	No check performed	
			LOGICAL BLOCK APPLICATION TAG	No check performed	
			LOGICAL BLOCK REFERENCE TAG	No check performed	
	No ^a	No protection information available to transmit to the Data-In Buffer or for checking			
100b ^b	Yes	Yes ^e	LOGICAL BLOCK GUARD	GRD_CHK = 1	LOGICAL BLOCK GUARD CHECK FAILED
				GRD_CHK = 0	No check performed
			LOGICAL BLOCK APPLICATION TAG	No check performed	
	LOGICAL BLOCK REFERENCE TAG	No check performed			
No ^a	No protection information available to transmit to the Data-In Buffer or for checking				

Table 58 – RDPROTECT field (part 3 of 3)

Code	Logical unit formatted with protection information	Shall device server transmit protection information?	Field in protection information ^h	Extended INQUIRY Data VPD page bit value ^g	If check fails ^{d f} , additional sense code
110b to 111b	Reserved				
<p>^a If the logical unit supports protection information (see 4.22) and has not been formatted with protection information, then the device server shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.</p> <p>^b If the logical unit does not support protection information, then the device server should terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.</p> <p>^c If the device server has knowledge of the contents of the LOGICAL BLOCK APPLICATION TAG field, then the device server shall check each logical block application tag. If the ATO bit in the Control mode page (see SPC-4) is set to one, then this knowledge is acquired from:</p> <ul style="list-style-type: none"> a) the EXPECTED LOGICAL BLOCK APPLICATION TAG field and the LOGICAL BLOCK APPLICATION TAG MASK field in the CDB, if a READ (32) command (see 5.14) is received by the device server; b) the Application Tag mode page (see 6.5.3), if a command other than READ (32) is received by the device server, and the ATMPE bit in the Control mode page (see SPC-4) is set to one; or c) a method not defined by this standard, if a command other than READ (32) is received by the device server, and the ATMPE bit is set to zero. <p>^d If the device server terminates the command with CHECK CONDITION status, then the device server shall set the sense key to ABORTED COMMAND.</p> <p>^e The device server shall transmit protection information to the Data-In Buffer.</p> <p>^f If multiple errors occur while the device server is processing the command, then the selection by the device server of which error to report is not defined by this standard.</p> <p>^g See the Extended INQUIRY Data VPD page (see SPC-4) for the definitions of the GRD_CHK bit, the APP_CHK bit, and the REF_CHK bit.</p> <p>^h If the device server detects:</p> <ul style="list-style-type: none"> a) a LOGICAL BLOCK APPLICATION TAG field set to FFFFh, and type 1 protection (see 4.22.2.3) or type 2 protection (see 4.22.2.4) is enabled; or b) a LOGICAL BLOCK APPLICATION TAG field set to FFFFh, the LOGICAL BLOCK REFERENCE TAG field set to FFFF_FFFFh, and type 3 protection (see 4.22.2.5) is enabled, <p>then the device server shall not check any protection information in the associated protection information interval.</p> <p>ⁱ If type 1 protection is enabled, then the device server shall check the logical block reference tag by comparing it to the lower four bytes of the LBA associated with the logical block. If type 2 protection or type 3 protection is enabled, and the device server has knowledge of the contents of the LOGICAL BLOCK REFERENCE TAG field, then the device server shall check each logical block reference tag. If type 2 protection is enabled, then this knowledge may be acquired through the EXPECTED INITIAL LOGICAL BLOCK REFERENCE TAG field in a READ (32) command (see 5.14). If type 3 protection is enabled, then the method for acquiring this knowledge is not defined by this standard.</p> <p>^j If the DPICZ bit in the Control mode page (see SPC-4) is set to one, then protection information shall not be checked.</p>					

A disable page out (DPO) bit set to zero specifies that the retention priority shall be determined by the RETENTION PRIORITY fields in the Caching mode page (see 6.5.5). A DPO bit set to one specifies that the device server shall assign the logical blocks accessed by this command the lowest retention priority for being fetched into or retained by the cache (see 4.15). A DPO bit set to one overrides any retention priority specified in the Caching mode page. All other aspects of the algorithm implementing the cache replacement strategy are not defined by this standard.

NOTE 7 - The DPO bit is used to control replacement of logical blocks in the cache when the application client has information on the future usage of the logical blocks. If the DPO bit is set to one, then the application client is specifying that the logical blocks accessed by the command are not likely to be accessed again in the near future and are not to be put in the cache nor retained by the cache. If the DPO bit is set to zero, then the application client is specifying that the logical blocks accessed by this command are likely to be accessed again in the near future.

A force unit access (FUA) bit set to one specifies that the device server shall read the logical blocks from:

- a) the non-volatile cache, if any; or
- b) the medium.

If the FUA bit is set to one and a volatile cache contains a more recent version of a logical block than the non-volatile cache, if any, or the medium, then, before reading the logical block, the device server shall write the logical block to:

- a) the non-volatile cache, if any; or
- b) the medium.

An FUA bit set to zero specifies that the device server may read the logical blocks from:

- a) the volatile cache, if any;
- b) the non-volatile cache, if any; or
- c) the medium.

If rebuild assist mode (see 4.20) is supported and not enabled, then the device server shall ignore the rebuild assist recovery control (RARC) bit. If rebuild assist mode is supported and enabled, then the RARC bit specifies that the device server shall perform read medium operations as defined in 4.20.3.2 and 4.20.3.3.

If the rebuild assist mode is not supported and the RARC bit is set to one, then the device server should terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

See the PRE-FETCH (10) command (see 5.8) for the definition of the LOGICAL BLOCK ADDRESS field.

See the PRE-FETCH (10) command and 4.23 for the definition of the GROUP NUMBER field.

The TRANSFER LENGTH field specifies the number of contiguous logical blocks of data that shall be read and transferred to the Data-In Buffer, starting with the logical block referenced by the LBA specified by the LOGICAL BLOCK ADDRESS field. A TRANSFER LENGTH field set to zero specifies that no logical blocks shall be read or transferred. This condition shall not be considered an error. Any other value specifies the number of logical blocks that shall be read and transferred. If the specified LBA and the specified transfer length exceed the capacity of the medium (see 4.5), then the device server shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to LOGICAL BLOCK ADDRESS OUT OF RANGE. The TRANSFER LENGTH field is constrained by the MAXIMUM TRANSFER LENGTH field in the Block Limits VPD page (see 6.6.3).

The CONTROL byte is defined in SAM-5.

5.12 READ (12) command

The READ (12) command (see table 59) requests that the device server perform the actions defined for the READ (10) command (see 5.11).

Migration from the READ (12) command to the READ (16) command is recommended for all implementations.

Table 59 – READ (12) command

Byte	Bit	7	6	5	4	3	2	1	0
0		OPERATION CODE (A8h)							
1		RDPROTECT			DPO	FUA	RARC	Obsolete	Obsolete
2	(MSB)	LOGICAL BLOCK ADDRESS							
...									
5									
6	(MSB)	TRANSFER LENGTH							
...									
9									
10		Restricted for MMC-6	Reserved		GROUP NUMBER				
11		CONTROL							

The OPERATION CODE field is defined in SPC-4 and shall be set to the value shown in table 59 for the READ (12) command.

The CONTROL byte is defined in SAM-5.

See the READ (10) command (see 5.11) for the definitions of the other fields in this command.

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC 14776-323:2017

5.13 READ (16) command

The READ (16) command (see table 60) requests that the device server perform the actions defined for the READ (10) command (see 5.11).

Table 60 – READ (16) command

Byte	Bit	7	6	5	4	3	2	1	0
0		OPERATION CODE (88h)							
1		RDPROTECT			DPO	FUA	RARC	Obsolete	Reserved
2	(MSB)	LOGICAL BLOCK ADDRESS							
...									
9									
10	(MSB)	TRANSFER LENGTH							
...									
13									
14		Restricted for MMC-6	Reserved		GROUP NUMBER				
15		CONTROL							

The OPERATION CODE field is defined in SPC-4 and shall be set to the value shown in table 60 for the READ (16) command.

The CONTROL byte is defined in SAM-5.

See the READ (10) command (see 5.11) for the definitions of the other fields in this command.

5.14 READ (32) command

The READ (32) command (see table 61) requests that the device server perform the actions defined for the READ (10) command (see 5.11).

The device server shall only process a READ (32) command if type 2 protection is enabled (see 4.22.2.4).

Table 61 – READ (32) command

Byte	Bit	7	6	5	4	3	2	1	0
0		OPERATION CODE (7Fh)							
1		CONTROL							
2		Reserved							
...									
5									
6		Reserved			GROUP NUMBER				
7		ADDITIONAL CDB LENGTH (18h)							
8	(MSB)	SERVICE ACTION (0009h)							
9									
10		RDPROTECT		DPO	FUA	RARC	Obsolete	Reserved	
11		Reserved							
12	(MSB)	LOGICAL BLOCK ADDRESS							
...									
19									
20	(MSB)	EXPECTED INITIAL LOGICAL BLOCK REFERENCE TAG							
...									
23									
24	(MSB)	EXPECTED LOGICAL BLOCK APPLICATION TAG							
25									
26	(MSB)	LOGICAL BLOCK APPLICATION TAG MASK							
27									
28	(MSB)	TRANSFER LENGTH							
...									
31									

The OPERATION CODE field, the ADDITIONAL CDB LENGTH field, and the SERVICE ACTION field are defined in SPC-4 and shall be set to the values shown in table 61 for the READ (32) command.

The CONTROL byte is defined in SAM-5.

See the READ (10) command (see 5.11) for the definitions of the GROUP NUMBER field, the RDPROTECT field, the DPO bit, the FUA bit, the RARC bit, the LOGICAL BLOCK ADDRESS field, and the TRANSFER LENGTH field.

If checking of the LOGICAL BLOCK REFERENCE TAG field is enabled (see table 58 in 5.11), then the EXPECTED INITIAL LOGICAL BLOCK REFERENCE TAG field contains the value of the LOGICAL BLOCK REFERENCE TAG field expected in the protection information of the first logical block accessed by the command instead of a value based on the LBA (see 4.22.3).

If the ATO bit is set to one in the Control mode page (see SPC-4), and checking of the LOGICAL BLOCK APPLICATION TAG field is enabled (see table 58 in 5.11), then the LOGICAL BLOCK APPLICATION TAG MASK field contains a value that is a bit mask for enabling the checking of the LOGICAL BLOCK APPLICATION TAG field in every instance of protection information for each logical block accessed by the command. A LOGICAL BLOCK APPLICATION TAG MASK field bit set to one enables the checking of the corresponding bit of the EXPECTED LOGICAL BLOCK APPLICATION TAG field with the corresponding bit of the LOGICAL BLOCK APPLICATION TAG field in every instance of protection information. A LOGICAL BLOCK APPLICATION TAG MASK field bit set to zero disables the checking of the corresponding bit of the EXPECTED LOGICAL BLOCK APPLICATION TAG field with the corresponding bit of the LOGICAL BLOCK APPLICATION TAG field in every instance of protection information.

If the ATO bit is set:

- a) to zero; or
- b) to one in the Control mode page (see SPC-4), and checking of the LOGICAL BLOCK APPLICATION TAG field is disabled (see table 58),

then the LOGICAL BLOCK APPLICATION TAG MASK field and the EXPECTED LOGICAL BLOCK APPLICATION TAG field shall be ignored

5.15 READ CAPACITY (10) command

5.15.1 READ CAPACITY (10) overview

The READ CAPACITY (10) command (see table 62) requests that the device server transfer eight bytes of parameter data describing the capacity and medium format of the direct access block device to the Data-In Buffer. This command may be processed as if it has a HEAD OF QUEUE task attribute (see 4.16). If the logical unit supports protection information (see 4.22) or logical block provisioning management (see 4.7), then the application client should use the READ CAPACITY (16) command (see 5.16) instead of the READ CAPACITY (10) command.

Migration from the READ CAPACITY (10) command to the READ CAPACITY (16) command is recommended for all implementations.

Table 62 – READ CAPACITY (10) command

Byte	Bit	7	6	5	4	3	2	1	0
0		OPERATION CODE (25h)							
1		Reserved							Obsolete
2		Obsolete							
...									
5		Obsolete							
6		Reserved							
7		Reserved							
8		Reserved							Obsolete
9		CONTROL							

The OPERATION CODE field is defined in SPC-4 and shall be set to the value shown in table 62 for the READ CAPACITY (10) command.

The CONTROL byte is defined in SAM-5.

5.15.2 READ CAPACITY (10) parameter data

The READ CAPACITY (10) parameter data is defined in table 63. Any time the READ CAPACITY (10) parameter data changes, the device server should establish a unit attention condition as described in .

Table 63 – READ CAPACITY (10) parameter data

Byte	Bit	7	6	5	4	3	2	1	0
0	(MSB)	RETURNED LOGICAL BLOCK ADDRESS							(LSB)
...									
3									
4	(MSB)	LOGICAL BLOCK LENGTH IN BYTES							(LSB)
...									
7									

The device server shall set the RETURNED LOGICAL BLOCK ADDRESS field to the lower of:

- a) the LBA of the last logical block on the direct access block device; or
- b) FFFF_FFFFh, if the LBA of the last logical block on the direct access block device is greater than the maximum value that is able to be specified in the RETURNED LOGICAL BLOCK ADDRESS field.

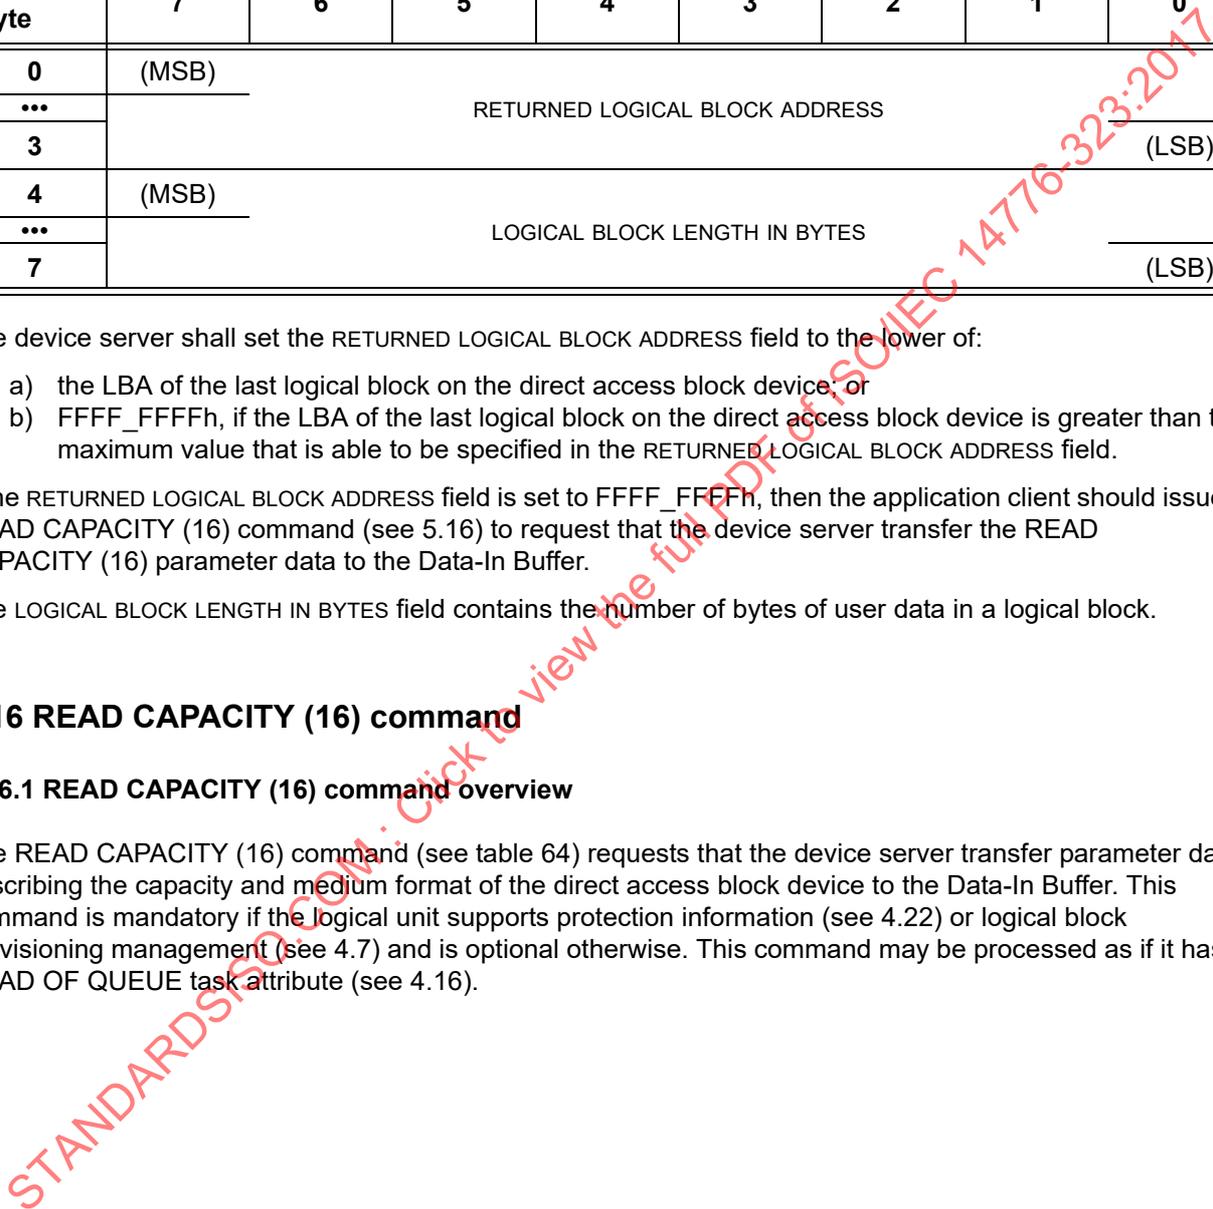
If the RETURNED LOGICAL BLOCK ADDRESS field is set to FFFF_FFFFh, then the application client should issue a READ CAPACITY (16) command (see 5.16) to request that the device server transfer the READ CAPACITY (16) parameter data to the Data-In Buffer.

The LOGICAL BLOCK LENGTH IN BYTES field contains the number of bytes of user data in a logical block.

5.16 READ CAPACITY (16) command

5.16.1 READ CAPACITY (16) command overview

The READ CAPACITY (16) command (see table 64) requests that the device server transfer parameter data describing the capacity and medium format of the direct access block device to the Data-In Buffer. This command is mandatory if the logical unit supports protection information (see 4.22) or logical block provisioning management (see 4.7) and is optional otherwise. This command may be processed as if it has a HEAD OF QUEUE task attribute (see 4.16).



This command uses the SERVICE ACTION IN (16) CDB format (see clause A.2).

Table 64 – READ CAPACITY (16) command

Bit Byte	7	6	5	4	3	2	1	0
0	OPERATION CODE (9Eh)							
1	Reserved			SERVICE ACTION (10h)				
2	Obsolete							
...								
9								
10	(MSB)	ALLOCATION LENGTH						(LSB)
...								
13								
14	Reserved						Obsolete	
15	CONTROL							

The OPERATION CODE field and SERVICE ACTION field are defined in SPC-4 and shall be set to the values shown in table 64 for the READ CAPACITY (16) command.

The ALLOCATION LENGTH field is defined in SPC-4.

The CONTROL byte is defined in SAM-5.

5.16.2 READ CAPACITY (16) parameter data

The READ CAPACITY (16) parameter data is defined in table 65. Any time the READ CAPACITY (16) parameter data changes, the device server should establish a unit attention condition as described in .

Table 65 – READ CAPACITY (16) parameter data

Byte	Bit	7	6	5	4	3	2	1	0
0	(MSB)	RETURNED LOGICAL BLOCK ADDRESS							(LSB)
...									
7									
8	(MSB)	LOGICAL BLOCK LENGTH IN BYTES							(LSB)
...									
11									
12	Reserved				P_TYPE			PROT_EN	
13	P_I_EXPONENT				LOGICAL BLOCKS PER PHYSICAL BLOCK EXPONENT				
14	LBPME	LBPRZ	(MSB)	LOWEST ALIGNED LOGICAL BLOCK ADDRESS					(LSB)
15									
16	Reserved								
...									
31									

The RETURNED LOGICAL BLOCK ADDRESS field and LOGICAL BLOCK LENGTH IN BYTES field of the READ CAPACITY (16) parameter data are defined in the READ CAPACITY (10) parameter data (see 5.15). The maximum value that shall be returned in the RETURNED LOGICAL BLOCK ADDRESS field is FFFF_FFFF_FFFF_FFFEh.

The protection type (P_TYPE) field and the protection enable (PROT_EN) bit (see table 66) indicate the logical unit's current type of protection.

Table 66 – P_TYPE field and PROT_EN bit

P_TYPE	PROT_EN	Description
n/a	0	The logical unit is formatted to type 0 protection (see 4.22.2.2).
000b	1	The logical unit is formatted to type 1 protection (see 4.22.2.3).
001b		The logical unit is formatted to type 2 protection (see 4.22.2.4).
010b		The logical unit is formatted to type 3 protection (see 4.22.2.5).
011b to 111b		Reserved

The P_I_EXPONENT field may be used to determine the number of protection information intervals placed within each logical block (see 5.3.2).

The number of protection information intervals is calculated as follows:

$$\text{number of protection information intervals} = 2^{(p_i \text{ exponent})}$$

where:

p_i exponent is the contents of the P_I_EXPONENT field.

The LOGICAL BLOCKS PER PHYSICAL BLOCK EXPONENT field is defined in table 67.

Table 67 – LOGICAL BLOCKS PER PHYSICAL BLOCK EXPONENT field

Code	Description
0	One or more physical blocks per logical block ^a
n > 0	2 ⁿ logical blocks per physical block
^a The number of physical blocks per logical block is not reported.	

A logical block provisioning management enabled (LBPME) bit set to one indicates that the logical unit implements logical block provisioning management (i.e., is resource provisioned or thin provisioned) (see 4.7.3). An LBPME bit set to zero indicates that the logical unit does not implement logical block provisioning management (e.g., is fully provisioned (see 4.7.2)).

A logical block provisioning read zeros (LBPRZ) bit set to one indicates that, for read commands specifying an unmapped LBA (see 4.7.4.5), the device server returns user data set to zero and protection information, if any, set to FFFF_FFFF_FFFF_FFFFh. An LBPRZ bit set to zero indicates that, for read commands specifying an unmapped LBA, the device server returns user data set to vendor specific data and protection information, if any, set to FFFF_FFFF_FFFF_FFFFh.

The LOWEST ALIGNED LOGICAL BLOCK ADDRESS field indicates the LBA of the first logical block that is located at the beginning of a physical block (see 4.6).

5.17 READ DEFECT DATA (10) command

5.17.1 READ DEFECT DATA (10) command overview

The READ DEFECT DATA (10) command (see table 68) requests that the device server transfer parameter data (see 5.17.2) containing a four-byte header, the PLIST, and/or the GLIST to the Data-In Buffer.

If the device server is unable to access a specified defect list due to a medium error, then the device server shall terminate the command with CHECK CONDITION status with the sense key set to MEDIUM ERROR and the additional sense code set to DEFECT LIST NOT FOUND.

If the device server is unable to access a specified defect list due to an error other than a medium error or because a specified defect list does not exist, then the device server shall either:

- 1) terminate the command with CHECK CONDITION status with the sense key set to NO SENSE and the additional sense code set to DEFECT LIST NOT FOUND; or
- 2) return only the READ DEFECT DATA parameter data header, with the DEFECT LIST LENGTH field set to zero.

Device servers may or may not return a defect list until after a successful completion of a FORMAT UNIT command (see 5.2).

Migration from the READ DEFECT DATA (10) command to the READ DEFECT DATA (12) command is recommended for all implementations.

Table 68 – READ DEFECT DATA (10) command

Byte	Bit	7	6	5	4	3	2	1	0
0		OPERATION CODE (37h)							
1		Reserved							
2		Reserved			REQ_PLIST	REQ_GLIST	DEFECT LIST FORMAT		
3		Reserved							
...		Reserved							
6		Reserved							
7	(MSB)	ALLOCATION LENGTH							
8		(LSB)							
9		CONTROL							

The OPERATION CODE field is defined in SPC-4 and shall be set to the value shown in table 68 for the READ DEFECT DATA (10) command.

Table 69 defines the request PLIST (REQ_PLIST) bit and the request GLIST (REQ_GLIST) bit.

Table 69 – REQ_PLIST bit and REQ_GLIST bit

REQ_PLIST	REQ_GLIST	Description
0	0	The device server shall return only the first four bytes of the READ DEFECT DATA parameter data (i.e., the parameter data header), with the DEFECT LIST LENGTH field set to zero.
	1	The device server shall return the READ DEFECT DATA parameter data header and include the GLIST, if any, in the defect list.
1	0	The device server shall return the READ DEFECT DATA parameter data header and include the PLIST, if any, in the defect list.
	1	The device server shall return the READ DEFECT DATA parameter data header and include the both the PLIST, if any, and the GLIST, if any, in the defect list. Whether the PLIST and GLIST are merged or not is vendor specific.

The DEFECT LIST FORMAT field specifies the address descriptor format type (see 6.2) that the device server should use for the defect list. A device server unable to return the requested address descriptor format shall return the address descriptors in their default format and indicate that format type in the DEFECT LIST FORMAT field in the READ DEFECT DATA parameter data header (see 5.17.2 and 5.18.2).

If the requested defect list format and the returned defect list format are not the same, then the device server shall transfer the defect data and then terminate the command with CHECK CONDITION status with the sense key set to RECOVERED ERROR and the additional sense code set to DEFECT LIST NOT FOUND.

The ALLOCATION LENGTH field is defined in SPC-4. If the length of the address descriptors that the device server has to report is greater than the maximum value that is able to be specified by the ALLOCATION LENGTH field, then the device server shall transfer no data and shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The CONTROL byte is defined in SAM-5.

STANDARDSISO.COM : Click to buy the PDF of ISO/IEC 14776-323:2017

5.17.2 READ DEFECT DATA (10) parameter data

The READ DEFECT DATA (10) parameter data (see table 70) contains a four-byte header, followed by zero or more address descriptors.

Table 70 – READ DEFECT DATA (10) parameter data

Byte	Bit	7	6	5	4	3	2	1	0
Parameter data header									
0	Reserved								
1	Reserved			PLISTV		GLISTV		DEFECT LIST FORMAT	
2	(MSB) _____ DEFECT LIST LENGTH (n - 3) _____								
3	(LSB)								
Defect list (if any)									
4	_____								
...	Address descriptor(s) (if any) _____								
n	_____								

A PLIST valid (PLISTV) bit set to zero indicates that the defect list does not contain the PLIST. A PLISTV bit set to one indicates that the defect list contains the PLIST.

A GLIST valid (GLISTV) bit set to zero indicates that the defect list does not contain the GLIST. A GLISTV bit set to one indicates that the defect list contains the GLIST.

The DEFECT LIST FORMAT field indicates the format of the address descriptors returned in the defect list. This field is defined in 6.2.

If the device server returns short block format address descriptors (see 6.2.2) or long block format address descriptors (see 6.2.5), then the address descriptors contain vendor specific values.

The DEFECT LIST LENGTH field indicates the length in bytes of the defect list. The DEFECT LIST LENGTH is equal to four or eight times the number of the address descriptors, depending on the format of the returned address descriptors (see 6.2).

The defect list contains address descriptors (see 6.2).

5.18 READ DEFECT DATA (12) command

5.18.1 READ DEFECT DATA (12) command overview

The READ DEFECT DATA (12) command (see table 71) requests that the device server transfer parameter data (see 5.18.2) containing a four-byte header, the PLIST, and/or the GLIST to the Data-In Buffer.

An application client determines the length of the defect list by sending a READ DEFECT DATA (12) command with an allocation length field set to eight and the address descriptor index field set to 0000_0000h. The device server returns the defect list header that contains the length of the defect list.

Table 71 – READ DEFECT DATA (12) command

Byte	Bit	7	6	5	4	3	2	1	0
0		OPERATION CODE (B7h)							
1		Reserved			REQ_PLIST	REQ_GLIST	DEFECT LIST FORMAT		
2	(MSB)	ADDRESS DESCRIPTOR INDEX							
...									
5									
6	(MSB)	ALLOCATION LENGTH							
...									
9									
10		Reserved							
11		CONTROL							

The OPERATION CODE field is defined in SPC-4 and shall be set to the value shown in table 71 for the READ DEFECT DATA (12) command.

See the READ DEFECT DATA (10) command (see 5.17) for the definitions of the REQ_PLIST bit, the REQ_GLIST bit, and the DEFECT LIST FORMAT field.

The ADDRESS DESCRIPTOR INDEX field specifies the index of the first address descriptor (see 6.2) in the defect list that the device server shall return. If the ADDRESS DESCRIPTOR INDEX field is set to:

- a) a value less than the number of available address descriptors, then the device server shall transfer a defect list beginning with the address descriptor that is at the ADDRESS DESCRIPTOR INDEX field value multiplied by the size of the address descriptor; or
- b) a value greater than or equal to the number of available address descriptors, then the device server shall return a zero length defect list.

The ALLOCATION LENGTH field is defined in SPC-4, however if the length of all the address descriptors that are available is greater than FFFF_FFFFh, then the device server shall transfer the length of address descriptors specified by the allocation length or the DEFECT LIST LENGTH field value plus eight, whichever is less, and complete the command with GOOD status.

The CONTROL byte is defined in SAM-5.

5.18.2 READ DEFECT DATA (12) parameter data

The READ DEFECT DATA (12) parameter data (see table 72) contains an eight-byte header, followed by zero or more address descriptors.

Table 72 – READ DEFECT DATA (12) parameter data

Byte	Bit	7	6	5	4	3	2	1	0
Parameter data header									
0	Reserved								
1	Reserved			PLISTV	GLISTV	DEFECT LIST FORMAT			
2	(MSB)	GENERATION CODE							(LSB)
3									
4	(MSB)	DEFECT LIST LENGTH (n - 7)							(LSB)
...									
7									
Defect list (if any)									
8									
...	Address descriptor(s) (if any)								
n									

The GENERATION CODE field is a two-byte counter that shall be incremented by one by the device server every time the defect list is changed. A GENERATION CODE field set to 0000h indicates the generation code is not supported. If the GENERATION CODE field is supported, then the GENERATION CODE field shall be initialized to at least 0001h at power on and the device server shall wrap this field to 0001h as the next increment after reaching its maximum value (i.e., FFFFh).

Application clients that use the GENERATION CODE field should read this field often enough to ensure that the contents of this field do not increment a multiple of 65 535 times between readings.

The DEFECT LIST LENGTH field indicates the length in bytes of address descriptors from the beginning address descriptor specified by the ADDRESS DESCRIPTOR INDEX field to the last address descriptor available to be returned. A value of FFFF_FFFFh in the DEFECT LIST LENGTH field indicates that more than FFFF_FFFEh bytes are available.

See the READ DEFECT DATA (10) command (see 5.17) for the definitions of the other fields in the READ DEFECT DATA (12) parameter data.

5.19 READ LONG (10) command

The READ LONG (10) command (see table 73) requests that the device server transfer data from a single logical block or physical block to the Data-In Buffer. The data transferred during the READ LONG (10) command is vendor specific, but shall include the following items recorded on the medium:

- a) if a logical block is being transferred, then:
 - A) user data or transformed user data for the logical block;
 - B) protection information or transformed protection information, if any, for the logical block; and
 - C) any additional information (e.g., ECC bytes) for all the physical blocks in the logical block;
 or
- b) if a physical block is being transferred, then:
 - A) user data or transformed user data for all the logical blocks in the physical block;
 - B) protection information or transformed protection information, if any, for all the logical blocks in the physical block; and
 - C) any additional information (e.g., ECC bytes).

If the additional information contain an ECC, then any other additional bytes that are correctable by ECC should be included (e.g., a data synchronization mark within the area covered by ECC). It is not required for the ECC bytes to be at the end of the logical block data. The ECC bytes should be in the same order as they are on the medium.

If a cache contains a more recent version of the specified logical block or physical block, then the device server shall write the logical block or physical block to the medium before reading it. The values in the Read-Write Error Recovery mode page (see 6.5.8) do not apply to this command. The device server may perform retries while processing this command.

This command uses the SERVICE ACTION IN (16) CDB format (see clause A.2).

Migration from the READ LONG (10) command to the READ LONG (16) command is recommended for all implementations.

Table 73 – READ LONG (10) command

Byte	Bit	7	6	5	4	3	2	1	0	
0		OPERATION CODE (3Eh)								
1		Reserved					PBLOCK	CORRCT	Obsolete	
2	(MSB)	LOGICAL BLOCK ADDRESS								
...										
5										(LSB)
6		Reserved								
7	(MSB)	BYTE TRANSFER LENGTH								
8										(LSB)
9										CONTROL

The OPERATION CODE field is defined in SPC-4 and shall be set to the value shown in table 73 for the READ LONG (10) command.

If there is more than one logical block per physical block (i.e., the LOGICAL BLOCKS PER PHYSICAL BLOCK EXPONENT field in the READ CAPACITY (16) parameter data (see 5.16.2) is set to a non-zero value), then:

- a) the device server shall support the physical block (PBLOCK) bit;

- b) a PBLOCK bit set to one specifies that the device server shall return the entire physical block containing the specified logical block; and
- c) a PBLOCK bit set to zero specifies that the device server shall return bytes representing only the specified logical block.

If there are one or more physical blocks per logical block (i.e., the LOGICAL BLOCKS PER PHYSICAL BLOCK EXPONENT field in the READ CAPACITY (16) parameter data (see 5.16.2) is set to zero), and the PBLOCK bit is set to one, then the device server shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

A correct (CORRCT) bit set to zero specifies that a logical block be read without any correction made by the device server. A CORRCT bit set to zero should result in the device server completing the command with GOOD status unless data is not transferred for some reason other than that the data is non-correctable. In this case the device server shall complete or terminate the command with the appropriate status and sense data. A CORRCT bit set to one specifies that the data be corrected by ECC before being transferred to the Data-In Buffer.

The LOGICAL BLOCK ADDRESS field specifies an LBA (see 4.5). If the specified LBA exceeds the capacity of the medium (see 4.5), then the device server shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to LOGICAL BLOCK ADDRESS OUT OF RANGE.

The BYTE TRANSFER LENGTH field specifies the number of bytes of data that shall be read from the specified logical block or physical block and transferred to the Data-In Buffer. If the BYTE TRANSFER LENGTH field is not set to zero and does not match the available data length, then the device server shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB. In the sense data (see 4.18 and SPC-4), the VALID bit shall be set to one and the ILI bit shall each be set to one, and the INFORMATION field shall be set to the difference (i.e., residue) of the requested byte transfer length minus the actual available data length in bytes. Negative values shall be indicated by two's complement notation.

A BYTE TRANSFER LENGTH field set to zero specifies that no bytes shall be read. This condition shall not be considered an error.

The CONTROL byte is defined in SAM-5.

5.20 READ LONG (16) command

The READ LONG (16) command (see table 74) requests that the device server perform the actions defined for the READ LONG (10) command (see 5.19).

Table 74 – READ LONG (16) command

Byte	Bit	7	6	5	4	3	2	1	0
0		OPERATION CODE (9Eh)							
1		Reserved			SERVICE ACTION (11h)				
2	(MSB)	LOGICAL BLOCK ADDRESS							
...									
9									
10		Reserved							
11		BYTE TRANSFER LENGTH							
12	(MSB)								
13									
14		Reserved						PBLOCK	CORRECT
15		CONTROL							

The OPERATION CODE field and SERVICE ACTION field are defined in SPC-4 and shall be set to the values shown in table 74 for the READ LONG (16) command.

The CONTROL byte is defined in SAM-5.

See the READ LONG (10) command (see 5.19) for the definitions of the other fields in this command.

5.21 REASSIGN BLOCKS command

5.21.1 REASSIGN BLOCKS command overview

The REASSIGN BLOCKS command (see table 75) requests that the device server perform a reassign operation on one or more LBAs (e.g., LBAs referencing logical blocks on which unrecovered read errors occurred) to another area on the medium set aside for this purpose and to add the physical blocks containing those logical blocks to the GLIST. This command shall not alter the contents of the PLIST (see 4.13).

The parameter list provided in the Data-Out Buffer contains a reassign LBA list that contains the LBAs of the logical blocks to be reassigned. The device server shall reassign the parts of the medium used for each logical block referenced by an LBA in the reassign LBA list. More than one physical block may be reassigned by each LBA. If the device server recovers logical block data from the original logical block, then the device server shall perform a write medium operation to that LBA using the recovered logical block data, which writes to the logical block referenced by the reassigned LBA.

The device server shall invalidate any of the specified LBAs that are in cache.

If the device server does not recover logical block data in a fully provisioned logical unit (see 4.7.2), then the device server shall:

- a) write vendor specific data as the user data; and
- b) write a default value of FFFF_FFFF_FFFF_FFFFh as the protection information, if enabled (see 4.22.2).

If the device server does not recover logical block data in a resource provisioned logical unit (see 4.7.3.2) or a thin provisioned logical unit (see 4.7.3.3), then the device server shall, for each specified LBA, either:

- a) unmap the specified LBA; or
- b) perform the following operations:
 - A) write vendor specific data as the user data; and
 - B) write a default value of FFFF_FFFF_FFFF_FFFFh as the protection information, if enabled.

The vendor specific data written as user data may contain remnants of the original logical block (e.g., partially or fully recovered user data).

The data in all other logical blocks on the medium shall be preserved.

Specifying an LBA to be reassigned that previously has been reassigned causes the device server to reassign that LBA again.

If the device server terminates the REASSIGN BLOCKS command with CHECK CONDITION status, and the sense data COMMAND-SPECIFIC INFORMATION field contains a valid LBA, then the application client should remove all LBAs from the reassign LBA list prior to the one returned in the COMMAND-SPECIFIC INFORMATION field. If the sense key is set to MEDIUM ERROR and the INFORMATION field contains the valid LBA, then the application client should insert that LBA into the reassign LBA list and reissue the REASSIGN BLOCKS command with the new reassign LBA list. Otherwise, the application client should perform any corrective action indicated by the sense data and then reissue the REASSIGN BLOCKS command with the new reassign LBA list.

Table 75 – REASSIGN BLOCKS command

Byte	Bit	7	6	5	4	3	2	1	0
0		OPERATION CODE (07h)							
1		Reserved						LONGLBA	LONGLIST
2		Reserved							
...									
4									
5		CONTROL							

The OPERATION CODE field is defined in SPC-4 and shall be set to the value shown in table 75 for the REASSIGN BLOCKS command.

A long LBA (LONGLBA) bit set to zero specifies that the reassign LBA list in the REASSIGN BLOCKS parameter list (see 5.21.2) contains four-byte LBAs. A LONGLBA bit set to one specifies that the reassign LBA list in the REASSIGN BLOCKS parameter list contains eight-byte LBAs.

A long list (LONGLIST) bit set to zero specifies the REASSIGN BLOCKS short parameter list header (see table 77) is used. A LONGLIST bit set to one specifies the REASSIGN BLOCKS long parameter list header (see table 78) is used.

The CONTROL byte is defined in SAM-5.

5.21.2 REASSIGN BLOCKS parameter list

The REASSIGN BLOCKS parameter list (see table 76) contains a four-byte parameter list header followed by a reassign LBA list containing one or more LBAs.

Table 76 – REASSIGN BLOCKS parameter list

Byte	Bit	7	6	5	4	3	2	1	0
0		Parameter list header (see table 77 or table 78)							
...									
3									
Reassign LBA list (if any)									
4		Reassign LBA [first] (see table 79 or table 80)							
...									
7 or 11									
n-4 or n-7		Reassign LBA [last] (see table 79 or table 80)							
...									
n									

The REASSIGN BLOCKS short parameter list header is defined in table 77.

Table 77 – REASSIGN BLOCKS short parameter list header

Byte	Bit	7	6	5	4	3	2	1	0
0		Reserved							
1									
2	(MSB)	REASSIGN LBA LENGTH							
3									
(LSB)									

The REASSIGN BLOCKS long parameter list header is defined in table 78.

Table 78 – REASSIGN BLOCKS long parameter list header

Byte	Bit	7	6	5	4	3	2	1	0
0	(MSB)	REASSIGN LBA LENGTH							
...									
3									
(LSB)									

The REASSIGN LBA LENGTH field specifies the total length in bytes of the reassign LBA list. The REASSIGN LBA LENGTH field does not include the parameter list header length and is equal to:

- a) four times the number of LBAs, if the `LONGLBA` bit in the REASSIGN BLOCKS command CDB bit is set to zero; or
- b) eight times the number of LBAs, if the `LONGLBA` bit is set to one.

The REASSIGN LBA LIST field contains a list of LBAs to be reassigned. The LBAs shall be sorted in ascending order.

If the LONGLBA bit is set to zero, then table 79 defines the format of the reassigned LBA.

Table 79 – Reassign LBA if the LONGLBA bit is set to zero

Byte	Bit	7	6	5	4	3	2	1	0	
0		(MSB)								
...		REASSIGN LOGICAL BLOCK ADDRESS								
3										(LSB)

The REASSIGN LOGICAL BLOCK ADDRESS field specifies an LBA to be reassigned.

If the LONGLBA bit is set to one, then table 80 defines the reassign LBA.

Table 80 – Reassign LBA if the LONGLBA bit is set to one

Byte	Bit	7	6	5	4	3	2	1	0	
0		(MSB)								
...		REASSIGN LOGICAL BLOCK ADDRESS								
7										(LSB)

The REASSIGN LOGICAL BLOCK ADDRESS field specifies an LBA to be reassigned.

If a specified LBA exceeds the capacity of the medium (see 4.5), then the device server shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code should be set to LOGICAL BLOCK ADDRESS OUT OF RANGE or may be set to INVALID FIELD IN PARAMETER LIST.

If the direct access block device has insufficient capacity to reassign all of the specified LBAs, then the device server shall terminate the command with CHECK CONDITION status with the sense key set to HARDWARE ERROR and the additional sense code set to NO DEFECT SPARE LOCATION AVAILABLE.

If the direct access block device is unable to complete a REASSIGN BLOCKS command without error, then the device server shall terminate the command with CHECK CONDITION status with the appropriate sense data (see 4.18 and SPC-4).

If one or more LBAs are not reassigned, then the device server shall report the first LBA not reassigned in the COMMAND-SPECIFIC INFORMATION field of the sense data (see SPC-4). If:

- a) information about the first LBA not reassigned is not available;
- b) all the LBAs have been reassigned; or
- c) the first LBA not reassigned does not fit in the COMMAND-SPECIFIC INFORMATION field, then the device server shall report the following value in the COMMAND-SPECIFIC INFORMATION field of the sense data (see SPC-4):
 - A) FFFF_FFFFh if fixed format sense data is being used; or
 - B) FFFF_FFFF_FFFF_FFFFh if descriptor format sense data is being used.

If the REASSIGN BLOCKS command failed due to an unexpected unrecovered read error that would cause the loss of data in a logical block not specified in the reassign LBA list, then the LBA of the logical block with the unrecovered read error is reported in the INFORMATION field of the sense data (see 4.18.1).

5.22 RECEIVE ROD TOKEN INFORMATION

5.22.1 RECEIVE ROD TOKEN INFORMATION overview

The RECEIVE ROD TOKEN INFORMATION command (see SPC-4) provides a method for an application client to receive information about the results of a previous or current block device ROD token operation. Table 81 shows the operations and a reference to the subclause where each topic is described.

Table 81 – RECEIVE ROD TOKEN INFORMATION reference

Command originating the operation	Command reference	RECEIVE ROD TOKEN INFORMATION returned parameter data reference
POPULATE TOKEN	5.7	5.22.2
WRITE USING TOKEN	5.46	5.22.3

5.22.2 RECEIVE ROD TOKEN INFORMATION parameter data for POPULATE TOKEN command

If a RECEIVE ROD TOKEN INFORMATION command (see SPC-4) specifies a list identifier that matches the list identifier specified in a previous POPULATE TOKEN command (see 5.7) received on the same I_T nexus, then table 82 shows the parameter data returned by the copy manager.

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC 14776-323:2017

Table 82 – RECEIVE ROD TOKEN INFORMATION parameter data for POPULATE TOKEN

Byte	Bit	7	6	5	4	3	2	1	0
0	(MSB)	AVAILABLE DATA (n - 3)							
...									
3									
4		Reserved			RESPONSE TO SERVICE ACTION (10h)				
5	Reserved	COPY OPERATION STATUS							
6	(MSB)	OPERATION COUNTER							
7									
8	(MSB)	ESTIMATED STATUS UPDATE DELAY							
...									
11									
12		EXTENDED COPY COMPLETION STATUS							
13		LENGTH OF THE SENSE DATA FIELD (m - 31)							
14		SENSE DATA LENGTH							
15		TRANSFER COUNT UNITS (F1h)							
16	(MSB)	TRANSFER COUNT							
...									
23									
24	(MSB)	SEGMENTS PROCESSED (0000h)							
25									
26		Reserved							
...									
31									
32		SENSE DATA (if any)							
...									
m									
m + 1	(MSB)	ROD TOKEN DESCRIPTOR LENGTH (n - (m + 4))							
...									
m + 4									
m + 5		Restricted (see SPC-4)							
m + 6		ROD TOKEN (if any)							
m + 7									
...									
n									

The AVAILABLE DATA field, the COPY OPERATION STATUS field, the OPERATION COUNTER field, the ESTIMATED STATUS UPDATE DELAY field, the EXTENDED COPY COMPLETION STATUS field, the LENGTH OF THE SENSE DATA FIELD field, SENSE DATA LENGTH field, the SENSE DATA field, and the ROD TOKEN field are defined in SPC-4.

The RESPONSE TO SERVICE ACTION field is defined in SPC-4 and shall be set to the value shown in table 82 in response to a RECEIVE ROD TOKEN INFORMATION command in which the LIST IDENTIFIER field specifies a POPULATE TOKEN command.

The TRANSFER COUNT UNITS field is defined in SPC-4 and shall be set to the value shown in table 82 in response to a RECEIVE ROD TOKEN INFORMATION command in which the LIST IDENTIFIER field specifies a POPULATE TOKEN command.

The TRANSFER COUNT field indicates the number of contiguous logical blocks represented by the ROD token that were read without error starting at the LBA specified in the first block device range descriptor and including the LBAs described in all complete block device range descriptors of the POPULATE TOKEN command to which this response applies.

If the value in the TRANSFER COUNT field is not equal to the sum of the contents of the NUMBER OF LOGICAL BLOCKS fields in all of the complete block device range descriptors of the POPULATE TOKEN command to which this response applies, then the COPY OPERATION STATUS field shall be set to 3h. Other values in the COPY OPERATION STATUS field are defined in SPC-4.

The SEGMENTS PROCESSED field is defined in SPC-4 and shall be set to the value shown in table 82 in response to a RECEIVE ROD TOKEN INFORMATION command in which the LIST IDENTIFIER field specifies a POPULATE TOKEN command.

The ROD TOKEN DESCRIPTOR LENGTH field is defined in SPC-4 and shall be set to the size of the ROD TOKEN field plus two in response to a RECEIVE ROD TOKEN INFORMATION command in which the LIST IDENTIFIER field specifies a POPULATE TOKEN command.

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC 14776-323:2017

5.22.3 The RECEIVE ROD TOKEN INFORMATION parameter data for the WRITE USING TOKEN command

If a RECEIVE ROD TOKEN INFORMATION command (see SPC-4) specifies a list identifier that matches the list identifier specified in a previous WRITE USING TOKEN command (see 5.46) received on the same I_T nexus, then table 83 shows the parameter data returned by the copy manager.

Table 83 – RECEIVE ROD TOKEN INFORMATION parameter data for WRITE USING TOKEN

Byte	Bit	7	6	5	4	3	2	1	0
0	(MSB)	AVAILABLE DATA (n - 3)							(LSB)
3		RESPONSE TO SERVICE ACTION (11h)							
4	Reserved	COPY OPERATION STATUS							
5	Reserved	OPERATION COUNTER							
6	(MSB)	ESTIMATED STATUS UPDATE DELAY							(LSB)
7		EXTENDED COPY COMPLETION STATUS							
8	(MSB)	LENGTH OF THE SENSE DATA FIELD ((n - 4) - 31)							
...		SENSE DATA LENGTH							
11		TRANSFER COUNT UNITS (F1h)							
12	(MSB)	TRANSFER COUNT							(LSB)
13		SEGMENTS PROCESSED (0000h)							(LSB)
14	(MSB)	Reserved							
15		SENSE DATA (if any)							
16		Reserved							
17		SENSE DATA (if any)							
18		Restricted (see SPC-4)							
19		Restricted (see SPC-4)							
20		Restricted (see SPC-4)							
21		Restricted (see SPC-4)							
22		Restricted (see SPC-4)							
23		Restricted (see SPC-4)							
24		Restricted (see SPC-4)							
25		Restricted (see SPC-4)							
26		Restricted (see SPC-4)							
27		Restricted (see SPC-4)							
28		Restricted (see SPC-4)							
29		Restricted (see SPC-4)							
30		Restricted (see SPC-4)							
31		Restricted (see SPC-4)							
32		Restricted (see SPC-4)							
...		Restricted (see SPC-4)							
n - 4		Restricted (see SPC-4)							
n - 3		Restricted (see SPC-4)							
n		Restricted (see SPC-4)							

The AVAILABLE DATA field, the COPY OPERATION STATUS field, the OPERATION COUNTER field, the ESTIMATED STATUS UPDATE DELAY field, the EXTENDED COPY COMPLETION STATUS field, the LENGTH OF THE SENSE DATA FIELD field, the SENSE DATA LENGTH field, and the SENSE DATA field are defined in SPC-4.

The RESPONSE TO SERVICE ACTION field is defined in SPC-4 and shall be set to the value shown in table 83 in response to a RECEIVE ROD TOKEN INFORMATION command in which the LIST IDENTIFIER field specifies a WRITE USING TOKEN command.

The TRANSFER COUNT UNITS field is defined in SPC-4 and shall be set to the value shown in table 83 in response to a RECEIVE ROD TOKEN INFORMATION command in which the LIST IDENTIFIER field specifies a WRITE USING TOKEN command.

The TRANSFER COUNT field indicates the number of contiguous logical blocks that were written without error starting with the LBA specified in the first block device range descriptor and including the LBAs specified in all block device range descriptors of the WRITE USING TOKEN command to which this response applies.

If the value in the TRANSFER COUNT field is not equal to the sum of the contents of the NUMBER OF LOGICAL BLOCKS fields in all of the complete block device range descriptors of the WRITE USING TOKEN command to which this response applies, then the COPY OPERATION STATUS field shall be set to 3h. Other values in the COPY OPERATION STATUS field are defined in SPC-4.

The SEGMENTS PROCESSED field is defined in SPC-4 and shall be set to the value shown in table 83 in response to a RECEIVE ROD TOKEN INFORMATION command in which the LIST IDENTIFIER field specifies a POPULATE TOKEN command.

5.23 REPORT REFERRALS command

5.23.1 REPORT REFERRALS command overview

The REPORT REFERRALS command (see table 84) requests that the device server transfer parameter data indicating the user data segment(s) on the logical unit and the SCSI target ports through which those user data segments may be accessed (see 4.28) to the Data-In Buffer. This command shall be supported by a logical unit that reports in the Extended INQUIRY Data VPD page (see SPC-4) that it supports referrals (i.e., the R_SUP bit set to one).

This command uses the SERVICE ACTION IN (16) CDB format (see clause A.2).

Table 84 – REPORT REFERRALS command

Byte	Bit	7	6	5	4	3	2	1	0
0		OPERATION CODE (9Eh)							
1		Reserved			SERVICE ACTION (13h)				
2	(MSB)	LOGICAL BLOCK ADDRESS							
9									
10	(MSB)	ALLOCATION LENGTH							
13									
14		Reserved							ONE_SEG
15		CONTROL							

The OPERATION CODE field and SERVICE ACTION field are defined in SPC-4 and shall be set to the values shown in table 84 for the REPORT REFERRALS command.

The LOGICAL BLOCK ADDRESS field specifies an LBA in the first user data segment that the device server shall report in the REPORT REFERRALS parameter data. If the specified LBA exceeds the capacity of the medium (see 4.5), then the device server shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to LOGICAL BLOCK ADDRESS OUT OF RANGE.

The ALLOCATION LENGTH field is defined in SPC-4.

A one segment (ONE_SEG) bit set to zero specifies that the device server shall return information on all user data segments starting with the user data segment that contains the LBA specified in the LOGICAL BLOCK ADDRESS field and ending with the user data segment that contains the last LBA of the logical unit. A ONE_SEG bit set to one specifies the device server shall only return information on the user data segment that contains the LBA specified in the LOGICAL BLOCK ADDRESS field.

The CONTROL byte is defined in SAM-5.

5.23.2 REPORT REFERRALS parameter data

The REPORT REFERRALS parameter data (see table 85) contains information indicating the user data segment(s) on the logical unit and the SCSI target port groups through which those user data segments may be accessed (see 4.28).

Table 85 – REPORT REFERRALS parameter data

Byte	Bit	7	6	5	4	3	2	1	0
0		Reserved							
1									
2	(MSB)	USER DATA SEGMENT REFERRAL DESCRIPTOR LENGTH (y - 3)							
3									
User data segment referral descriptor list									
4		User data segment referral descriptor [first] (if any)							
...									
4 + n									
⋮									
y - m		User data segment referral descriptor [last] (if any)							
...									
y									

The USER DATA SEGMENT REFERRAL DESCRIPTOR LENGTH field indicates the number of bytes that follow in the REPORT REFERRALS parameter data.

The user data segment referral descriptor (see table 16) is defined in the user data segment referral sense data descriptor (see 4.18.4).

STANDARDSPRO.COM: Click to view the full PDF of ISO/IEC 14776-323:2017

5.24 SANITIZE command

5.24.1 SANITIZE command overview

The SANITIZE command (see table 86) requests that the device server perform a sanitize operation (see 4.11). This device server shall process this command as if it has a HEAD OF QUEUE task attribute (see 4.16).

Table 86 – SANITIZE command

Byte	Bit	7	6	5	4	3	2	1	0	
0		OPERATION CODE (48h)								
1		IMMED	Reserved	AUSE	SERVICE ACTION					
2		Reserved								
...										
6										
7		(MSB)	PARAMETER LIST LENGTH							
8									LSB)	
9		CONTROL								

The OPERATION CODE field is defined in SPC-4 and shall be set to the value shown in table 86 for the SANITIZE command.

If the immediate (IMMED) bit is set to zero, then the device server shall return status after the sanitize operation is completed. If the IMMED bit set to one, then the device server shall return status as soon as the CDB and parameter data, if any, have been validated. The REQUEST SENSE command may be used to poll for progress of the sanitize operation regardless of the value of the IMMED bit.

If the allow unrestricted sanitize exit (AUSE) bit is set to one, and the specified sanitize operation fails, then the device server shall process a subsequent EXIT FAILURE MODE service action as if the previous sanitize operation had completed without error (see 4.11.3)

If:

- a) the AUSE bit is set to zero in the SANITIZE command that requested a sanitize operation;
- b) the specified sanitize operation completes with an error; and
- c) a subsequent SANITIZE command with the EXIT FAILURE MODE service action is received,

then the device sever shall terminate that subsequent SANITIZE command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The SERVICE ACTION field is defined in 5.24.2.

The PARAMETER LIST LENGTH field specifies the length in bytes of the parameter data that is available to be transferred from the Data-Out Buffer. A PARAMETER LIST LENGTH field set to zero specifies that no data shall be transferred.

The CONTROL byte is defined in SAM-5.

5.24.2 SANITIZE command service actions

5.24.2.1 SANITIZE command service actions overview

The SANITIZE command service actions are defined in table 87. At least one service action shall be supported if the SANITIZE command is supported. If the service action specified in the CDB is not supported, then the device server shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

If deferred microcode has been saved and not activated (see SPC-4), then the device server shall terminate this command with CHECK CONDITION status with the sense key set to NOT READY and the additional sense code set to LOGICAL UNIT NOT READY, MICROCODE ACTIVATION REQUIRED.

Table 87 – SANITIZE service action codes

Code	Name	Description	PARAMETER LIST LENGTH requirement ^a	Reference
01h	OVERWRITE	Perform a sanitize overwrite operation	Set to > 0004h and < (logical block length + 5)	5.24.2.2
02h	BLOCK ERASE	Perform a sanitize block erase operation	Set to 0000h	5.24.2.3
03h	CRYPTOGRAPHIC ERASE	Perform a sanitize cryptographic erase operation	Set to 0000h	5.24.2.4
1Fh	EXIT FAILURE MODE	Exit the sanitize failure mode	Set to 0000h	5.24.2.5
all others	Reserved			
^a If the requirement is not met, then the SANITIZE command is terminated with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.				

5.24.2.2 OVERWRITE service action

The OVERWRITE service action (see table 87) requests that the device server perform a sanitize overwrite operation (see 4.11).

The parameter list format for the OVERWRITE service action is defined in table 88.

Table 88 – OVERWRITE service action parameter list

Byte	Bit	7	6	5	4	3	2	1	0
0		INVERT	TEST		OVERWRITE COUNT				
1		Reserved							
2	(MSB)	INITIALIZATION PATTERN LENGTH (n - 3)							
3		LSB)							
4		INITIALIZATION PATTERN							
...									
n									

If the INVERT bit is set to zero, then on each overwrite pass:

- a) the user data shall be written as specified in the INITIALIZATION PATTERN field; and
- b) the protection information, if any, shall be set to FFFF_FFFF_FFFF_FFFFh.

If the INVERT bit is set to one, then the user data and protection information bytes, if any, shall be inverted (i.e., each bit XORed with one) between consecutive overwrite passes.

The TEST field is defined in table 89.

Table 89 – TEST field

Code	Description
00b	Shall not cause any changes in the defined behavior of the SANITIZE command.
01b to 11b	Vendor specific ^a
^a Setting the TEST field to one of these values may adversely affect security properties of the OVERWRITE service action.	

The OVERWRITE COUNT field specifies the number of overwrite passes to be performed. The value of 00h is reserved.

The INITIALIZATION PATTERN LENGTH field specifies the length in bytes of the INITIALIZATION PATTERN field. If the INITIALIZATION PATTERN LENGTH field is set to zero or a value greater than the logical block length, then the device server shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST.

The INITIALIZATION PATTERN field specifies the data pattern to be used to write the user data. This data pattern is repeated as necessary to fill each logical block. For each logical block, the first byte of the user data shall begin with the first byte of the initialization pattern.

If the INVERT bit is set to one and:

- a) the OVERWRITE COUNT field is set to an even number, then the pattern used for the first write pass shall consist of:
 - A) the user data set to the inversion of the INITIALIZATION PATTERN data; and
 - B) the protection information, if any, set to 0000_0000_0000_0000h;
 or
- b) the OVERWRITE COUNT field is set to an odd number, then the pattern used for the first write pass shall consist of:
 - A) the user data set to the INITIALIZATION PATTERN data; and
 - B) the protection information, if any, set to FFFF_FFFF_FFFF_FFFFh.

After a sanitize overwrite operation completes without error:

- a) the device server completes read commands for which no other error occurs during processing with GOOD status and read medium operations return the data written by the sanitize overwrite operation; and
- b) protection information, if any, shall be set to FFFF_FFFF_FFFF_FFFFh in all logical blocks on the medium.

5.24.2.3 BLOCK ERASE service action

The BLOCK ERASE service action (see table 87) requests that the device server perform a sanitize block erase operation (see 4.11).

After a sanitize block erase operation completes without error:

- a) the device server may terminate commands that request read operations specifying mapped LBAs (see 4.7.1) based on the setting of the WABEREQ field in the Block Device Characteristics VPD page (see 6.6.2); and
- b) if the logical unit is formatted with protection information, then:
 - A) the protection information for each mapped LBA may be indeterminate; and
 - B) if the device server terminates a command that requests read operations specifying mapped LBAs as a result of a protection information error, then the device server shall terminate that command with CHECK CONDITION status with the sense key set to ABORTED COMMAND and the appropriate additional sense code for the condition (e.g., for READ commands, the additional sense code defined in table 58).

5.24.2.4 CRYPTOGRAPHIC ERASE service action

The CRYPTOGRAPHIC ERASE service action (see table 87) requests that the device server perform a sanitize cryptographic erase operation (see 4.11).

After a sanitize cryptographic erase operation completes without error:

- a) the device server may terminate commands that request read operations specifying mapped LBAs (see 4.7.1) based on the setting of the WACEREQ field in the Block Device Characteristics VPD page (see 6.6.2); and
- b) if the logical unit is formatted with protection information, then:
 - A) the protection information for each mapped LBA may be indeterminate; and
 - B) if the device server terminates a command that requests read operations specifying mapped LBAs as a result of a protection information error, then the device server shall terminate that command with CHECK CONDITION status with the sense key set to ABORTED COMMAND and the appropriate additional sense code for the condition (e.g., for READ commands, the additional sense code defined in table 58).

5.24.2.5 EXIT FAILURE MODE service action

The EXIT FAILURE MODE service action (see table 87) requests that the device server complete a sanitize operation which completed with an error as if the sanitize operation completed without an error (see 4.11). If the most recent sanitize operation, if any, has completed without error, then the EXIT FAILURE MODE service action completes without error.

After successful completion of a SANITIZE command with the EXIT FAILURE MODE service action:

- a) if any LBA is mapped (see 4.7.1), and the logical unit is formatted with protection information, then:
 - A) the protection information for each mapped LBA may be indeterminate; and
 - B) if the device server terminates a command that requests read operations specifying mapped LBAs as a result of a protection information error, then the device server shall terminate that command with CHECK CONDITION status with the sense key set to ABORTED COMMAND and the appropriate additional sense code for the condition (e.g., for READ commands, the additional sense code defined in table 58).
- and
- b) the device server should complete reads to unmapped LBAs without error (see 4.7.4.7.1 and 4.7.4.8.1).

5.25 START STOP UNIT command

The START STOP UNIT command (see table 90) requests that the device server change the power condition of the logical unit (see 4.21) or load or eject the medium. This includes specifying that the device server enable or disable the direct access block device for medium access operations by controlling power conditions and timers.

The device server shall handle any deferred microcode as specified in 4.26.

Table 90 – START STOP UNIT command

Byte	Bit	7	6	5	4	3	2	1	0
0	OPERATION CODE (1Bh)								
1	Reserved								IMMED
2	Reserved								
3	Reserved				POWER CONDITION MODIFIER				
4	POWER CONDITION				Reserved	NO_FLUSH	LOEJ	START	
5	CONTROL								

The OPERATION CODE field is defined in SPC-4 and shall be set to the value shown in table 90 for the START STOP UNIT command.

If the immediate (IMMED) bit is set to zero, then the device server shall return status after the operation is completed. If the IMMED bit is set to one, then the device server shall return status as soon as the CDB has been validated.

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC 14776-323:2017

The combinations of values in the POWER CONDITION field and the POWER CONDITION MODIFIER field are defined in table 91. If the POWER CONDITION field is supported and is set to a value other than 0h, then the device server shall ignore the START bit and the LOEJ bit.

Table 91 – POWER CONDITION and POWER CONDITION MODIFIER field (part 1 of 2)

POWER CONDITION value	POWER CONDITION name	POWER CONDITION MODIFIER value	Device server action
0h	START_VALID	0h	Process the START bit and the LOEJ bit.
1h	ACTIVE	0h	Cause the logical unit to transition to the active power condition. ^a
2h	IDLE	0h	Cause the logical unit to transition to the idle_a power condition. ^{a b}
		1h	Cause the logical unit to transition to the idle_b power condition. ^{a b c}
		2h	Cause the logical unit to transition to the idle_c power condition. ^{a b d}
3h	STANDBY	0h	Cause the logical unit to transition to the standby_z power condition. ^{a b}
		1h	Cause the logical unit to transition to the standby_y power condition. ^{a b}
5h	Obsolete	0h to Fh	Obsolete
7h	LU_CONTROL	0h	Initialize and start all of the idle condition timers that are enabled (see SPC-4), and initialize and start all of the standby condition timers that are enabled (see SPC-4).

^a Process the following actions:

- 1) The device server shall comply with any SCSI transport protocol specific power condition transition restrictions (e.g., the NOTIFY (ENABLE SPINUP) requirement (see SPL-2));
- 2) the logical unit shall transition to the specified power condition; and
- 3) the device server shall disable all of the idle condition timers that are enabled (see SPC-4) and disable all of the standby condition timers that are enabled (see SPC-4) until another START STOP UNIT command is processed that returns control of the power condition to the logical unit or a logical unit reset occurs.

^b If a timer for a background scan operation expires, or a device specific event occurs, then the logical unit shall not leave this power condition to perform the background function associated with the timer or event.

^c The device server shall cause the direct access block device to increase its tolerance of external physical forces (e.g., causes a device that has movable read/write heads to move those heads to a safe position).

^d The device server shall cause the direct access block device to increase its tolerance of external physical forces and reduce its power consumption to use less power than when the logical unit is in the idle_b power condition (e.g., cause a device that has rotating medium to rotate the medium at a lower revolutions per minute).

^e If the specified timer is supported and enabled, then the device server shall:

- 1) force the specified timer to expire, which may cause the logical unit to transition to the specified power condition;
- 2) initialize and start all of the idle condition timers that are enabled (see SPC-4); and
- 3) initialize and start all of the standby condition timers that are enabled (see SPC-4),

otherwise the device server shall terminate the START STOP UNIT command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

Table 91 – POWER CONDITION and POWER CONDITION MODIFIER field (part 2 of 2)

POWER CONDITION value	POWER CONDITION name	POWER CONDITION MODIFIER value	Device server action
Ah	FORCE_IDLE_0	0h	Force the idle_a condition timer to be initialized to zero. ^e
		1h	Force the idle_b condition timer to be initialized to zero. ^e
		2h	Force the idle_c condition timer to be initialized to zero. ^e
Bh	FORCE_STANDBY_0	0h	Force the standby_z condition timer to be initialized to zero. ^e
		1h	Force the standby_y condition timer to be initialized to zero. ^e
All other combinations			Reserved

^a Process the following actions:

- 1) The device server shall comply with any SCSI transport protocol specific power condition transition restrictions (e.g., the NOTIFY (ENABLE SPINUP) requirement (see SPL-2));
- 2) the logical unit shall transition to the specified power condition; and
- 3) the device server shall disable all of the idle condition timers that are enabled (see SPC-4) and disable all of the standby condition timers that are enabled (see SPC-4) until another START STOP UNIT command is processed that returns control of the power condition to the logical unit or a logical unit reset occurs.

^b If a timer for a background scan operation expires, or a device specific event occurs, then the logical unit shall not leave this power condition to perform the background function associated with the timer or event.

^c The device server shall cause the direct access block device to increase its tolerance of external physical forces (e.g., causes a device that has movable read/write heads to move those heads to a safe position).

^d The device server shall cause the direct access block device to increase its tolerance of external physical forces and reduce its power consumption to use less power than when the logical unit is in the idle_b power condition (e.g., cause a device that has rotating medium to rotate the medium at a lower revolutions per minute).

^e If the specified timer is supported and enabled, then the device server shall:

- 1) force the specified timer to expire, which may cause the logical unit to transition to the specified power condition;
- 2) initialize and start all of the idle condition timers that are enabled (see SPC-4); and
- 3) initialize and start all of the standby condition timers that are enabled (see SPC-4), otherwise the device server shall terminate the START STOP UNIT command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

If the START STOP UNIT command specifies a power condition that conflicts with an operation in progress then, after the START STOP UNIT command completes with GOOD status, the logical unit may not be in the power condition that was requested by the command.

It is not an error to specify that the logical unit transition to its current power condition.

If no START STOP UNIT command is being processed by the device server, then the device server shall process any received START STOP UNIT command.

If a START STOP UNIT command is being processed by the device server and a subsequent START STOP UNIT command for which the CDB is validated requests that the logical unit change to the same power condition that was specified by the START STOP UNIT command being processed, then the device server shall process the subsequent command.

If the NO_FLUSH bit is set to zero, then the device server shall write all logical block data in cache that is newer than logical block data on the medium to the medium (e.g., as if in response to a SYNCHRONIZE CACHE command (see 5.26 and 5.27) with the LOGICAL BLOCK ADDRESS field set to zero and the NUMBER OF LOGICAL BLOCKS field set to zero) prior to entering into any power condition that prevents accessing the medium (e.g.,

before the rotating medium spindle motor is stopped during transition to the stopped power condition). If the NO_FLUSH bit is set to one, then the device server should not write any cached logical blocks to the medium prior to entering into any power condition that prevents accessing the medium.

If the load eject (LOEJ) bit is set to zero and the POWER CONDITION field is set to zero, then the logical unit shall take no action regarding loading or ejecting the medium. If the LOEJ bit is set to one and the POWER CONDITION field is set to zero, then the logical unit shall unload the medium if the START bit is set to zero. If the LOEJ bit is set to one, the POWER CONDITION field is set to zero, and the START bit is set to one, then the logical unit shall load the medium.

If the START bit is set to zero and the POWER CONDITION field is set to zero, then the device server shall:

- a) cause the logical unit to transition to the stopped power condition;
- b) stop any idle condition timer that is enabled (see SPC-4); and
- c) stop any standby condition timer that is enabled (see SPC-4).

If the START bit set to one and the POWER CONDITION field is set to zero, then the device server shall:

- 1) comply with requirements defined in SCSI transport protocol standards (e.g., the NOTIFY (ENABLE SPINUP) requirement (see SPL-2));
- 2) cause the logical unit to transition to the active power condition;
- 3) initialize and start any idle condition timer that is enabled; and
- 4) initialize and start any standby condition timer that is enabled.

The CONTROL byte is defined in SAM-5.

5.26 SYNCHRONIZE CACHE (10) command

The SYNCHRONIZE CACHE (10) command (see table 92) requests that, for each logical block whose logical block data is in the volatile cache and has not already been written to the non-volatile cache, if any, or the medium, the device server either:

- a) perform a write medium operation to the LBA using the logical block data in volatile cache; or
- b) write the logical block to the non-volatile cache, if any.

Migration from the SYNCHRONIZE CACHE (10) command to the SYNCHRONIZE CACHE (16) command is recommended for all implementations.

Table 92 – SYNCHRONIZE CACHE (10) command

Byte	Bit	7	6	5	4	3	2	1	0	
0		OPERATION CODE (35h)								
1		Reserved				Obsolete		IMMED	Obsolete	
2	(MSB)	LOGICAL BLOCK ADDRESS								
...										
5		(LSB)								
6		Reserved			GROUP NUMBER					
7	(MSB)	NUMBER OF LOGICAL BLOCKS								
8										(LSB)
9		CONTROL								

The OPERATION CODE field is defined in SPC-4 and shall be set to the value shown in table 92 for the SYNCHRONIZE CACHE (10) command.

See the PRE-FETCH (10) command (see 5.8) for the definition of the LOGICAL BLOCK ADDRESS field.

See the PRE-FETCH (10) command (see 5.8) and 4.23 for the definition of the GROUP NUMBER field.

An immediate (IMMED) bit set to zero specifies that the device server shall not return status until the synchronize cache operation has been completed. An IMMED bit set to one specifies that the device server shall return status as soon as the CDB has been validated. If the IMMED bit is set to one, and the device server does not support the IMMED bit, then the device server shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

If the IMMED bit is set to one and the synchronize cache operation has not completed, then the SYNC_PROG field in the Caching mode page (see 6.5.5) defines the device server behavior while the synchronize cache command is being processed.

A NUMBER OF LOGICAL BLOCKS field set to a non-zero value specifies the number of logical blocks that shall be synchronized, starting with the logical block referenced by the LBA specified by the LOGICAL BLOCK ADDRESS field. A NUMBER OF LOGICAL BLOCKS field set to zero specifies that all logical blocks starting with the one referenced by the LBA specified in the LOGICAL BLOCK ADDRESS field to the last logical block on the medium shall be synchronized. If the specified LBA and the specified number of logical blocks exceed the capacity of the medium (see 4.5), then the device server shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to LOGICAL BLOCK ADDRESS OUT OF RANGE.

A logical block within the range that is not in cache is not considered an error.

The CONTROL byte is defined in SAM-5.

5.27 SYNCHRONIZE CACHE (16) command

The SYNCHRONIZE CACHE (16) command (see table 93) requests that the device server perform the actions defined for the SYNCHRONIZE CACHE (10) command (see 5.26).

Table 93 – SYNCHRONIZE CACHE (16) command

Byte	Bit	7	6	5	4	3	2	1	0	
0		OPERATION CODE (91h)								
1		Reserved				Obsolete		IMMED	Reserved	
2	(MSB)	LOGICAL BLOCK ADDRESS								
...										
9										(LSB)
10	(MSB)	NUMBER OF LOGICAL BLOCKS								
...										
13										(LSB)
14		Reserved				GROUP NUMBER				
15		CONTROL								

The OPERATION CODE field is defined in SPC-4 and shall be set to the value shown in table 93 for the SYNCHRONIZE CACHE (16) command.

See the SYNCHRONIZE CACHE (10) command (see 5.26) for the definitions of the other fields in this command.

5.28 UNMAP command

5.28.1 UNMAP command overview

The UNMAP command (see table 94) requests that the device server cause one or more LBAs to be unmapped (see 4.7.3).

Table 94 – UNMAP command

Byte	Bit	7	6	5	4	3	2	1	0
0		OPERATION CODE (42h)							
1		Reserved							
2		Reserved							
...									
5									
6		Reserved			GROUP NUMBER				
7	(MSB)	PARAMETER LIST LENGTH							
8									
9		CONTROL							

The OPERATION CODE field is defined in SPC-4 and shall be set to the value shown in table 94 for the UNMAP command.

For a thin provisioned logical unit (see 4.7.3.3):

- an ANCHOR bit set to zero specifies that any LBA on which an unmap operation (see 4.7.3.4) is performed shall either become deallocated or anchored and should become deallocated; and
- an ANCHOR bit set to one specifies that any LBA on which an unmap operation is performed shall become anchored.

For a resource provisioned logical unit (see 4.7.3.2), any LBA on which an unmap operation is performed shall become anchored (i.e., the command is processed as if the ANCHOR bit is set to one).

If the ANCHOR bit is set to one, and the ANC_SUP bit in the Logical Block Provisioning VPD page (see 6.6.4) is set to zero, then the device server shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

See the PRE-FETCH (10) command (see 5.8) and 4.23 for the definition of the GROUP NUMBER field.

The PARAMETER LIST LENGTH field specifies the length in bytes of the UNMAP parameter list that is available to be transferred from the Data-Out Buffer. If the parameter list length is greater than zero and less than eight, then the device server shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to PARAMETER LIST LENGTH ERROR. A PARAMETER LIST LENGTH set to zero specifies that no data shall be transferred.

The CONTROL byte is defined in SAM-5.

5.28.2 UNMAP parameter list

The UNMAP parameter list (see table 95) contains an UNMAP parameter list header and block descriptors for LBA extents for which unmap requests (see 4.7.3.4.1) are to be processed by the device server. The LBAs specified in the block descriptors may contain overlapping LBA extents, and may be in any order.

If the ANCHOR bit in the CDB is set to zero and the logical unit is thin provisioned (see 4.7.3.3), then the logical block provisioning state for each specified LBA:

- a) should become deallocated;
- b) may become anchored; or
- c) may remain unchanged.

If:

- a) the ANCHOR bit in the CDB is set to one and the ANC_SUP bit in the Logical Block Provisioning VPD page (see 6.6.4) is set to one; or
- b) the logical unit is resource provisioned,

then for each specified LBA:

- a) if the LBA is mapped, then the LBA shall not become deallocated and:
 - A) that LBA should become anchored (see 4.7.4.6.3); or
 - B) the logical block provisioning state of that LBA remains unchanged;
- b) if the LBA is deallocated, then that LBA shall become anchored (see 4.7.4.7.3). If a lack of LBA mapping resources prevents the LBA from becoming anchored, then the device server shall terminate the command as described in 4.7.3.6; or
- c) if the LBA is anchored, then that LBA shall remain anchored.

Table 95 – UNMAP parameter list

Byte	Bit	7	6	5	4	3	2	1	0	
0	(MSB)	UNMAP DATA LENGTH (n - 1)								
1									(LSB)	
2	(MSB)	UNMAP BLOCK DESCRIPTOR DATA LENGTH (n - 7)								
3									(LSB)	
4		Reserved								
...										
7										
UNMAP block descriptor list (if any)										
8		UNMAP block descriptor [first] (see table 96) (if any)								
...										
23										
		⋮								
n - 15		UNMAP block descriptor [last] (see table 96) (if any)								
...										
n										

The UNMAP DATA LENGTH field specifies the length in bytes of the following data that is available to be transferred from the Data-Out Buffer. The unmap data length does not include the number of bytes in the UNMAP DATA LENGTH field.

The UNMAP BLOCK DESCRIPTOR DATA LENGTH field specifies the length in bytes of the UNMAP block descriptors that are available to be transferred from the Data-Out Buffer. The unmap block descriptor data length should be a multiple of 16. If the unmap block descriptor data length is not a multiple of 16, then the last unmap block descriptor is incomplete and shall be ignored. If the UNMAP BLOCK DESCRIPTOR DATA LENGTH is set to zero, then no unmap block descriptors are included in the UNMAP parameter list. This condition shall not be considered an error.

If any UNMAP block descriptors in the UNMAP block descriptor list are truncated due to the parameter list length in the CDB, then that UNMAP block descriptor shall be ignored.

Table 96 defines the UNMAP block descriptor.

Table 96 – UNMAP block descriptor

Byte	Bit	7	6	5	4	3	2	1	0
0	(MSB)	UNMAP LOGICAL BLOCK ADDRESS							
...									
7									
8	(MSB)	NUMBER OF LOGICAL BLOCKS							
...									
11									
12		Reserved							
...									
15									

The UNMAP LOGICAL BLOCK ADDRESS field specifies the first LBA that is requested to be unmapped (see 4.7.3.4.1) for this UNMAP block descriptor.

The NUMBER OF LOGICAL BLOCKS field specifies the number of LBAs that are requested to be unmapped beginning with the LBA specified by the UNMAP LOGICAL BLOCK ADDRESS field.

If the NUMBER OF LOGICAL BLOCKS is set to zero, then no LBAs shall be unmapped for this UNMAP block descriptor. This condition shall not be considered an error.

If the specified LBA and the specified number of logical blocks exceeds the capacity of the medium (see 4.5), then the device server shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to LOGICAL BLOCK ADDRESS OUT OF RANGE.

If the total number of logical blocks specified in the UNMAP block descriptor data exceeds the value indicated in the MAXIMUM UNMAP LBA COUNT field in the Block Limits VPD page (see 6.6.3) or if the number of UNMAP block descriptors exceeds the value of the MAXIMUM UNMAP BLOCK DESCRIPTOR COUNT field in the Block Limits VPD page, then the device server shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST.

5.29 VERIFY (10) command

The VERIFY (10) command (see table 98) requests that the device server:

- 1) perform verify operations from the specified LBAs; and
- 2) if specified, perform a compare operation on:
 - A) the logical block data transferred from the Data-Out Buffer; and
 - B) the logical block data from the verify operations.

The device server may process the LBAs in any order but shall perform this sequence in the specified order for a given LBA.

The application client uses the BYTCHK field in the CDB to specify the contents of the Data-Out Buffer as shown in table 97.

Table 97 – Data-Out Buffer contents for the VERIFY (10) command

BYTCHK field	Data-Out Buffer contents
00b	Not used
01b	Logical block data for the number of logical blocks specified in the VERIFICATION LENGTH field
10b	Not defined
11b	Logical block data for a single logical block

Migration from the VERIFY (10) command to the VERIFY (16) command is recommended for all implementations.

Table 98 – VERIFY (10) command

Byte	Bit	7	6	5	4	3	2	1	0
0		OPERATION CODE (2Fh)							
1		VRPROTECT			DPO	Reserved	BYTCHK		Obsolete
2		(MSB)							
...		LOGICAL BLOCK ADDRESS							
5		(LSB)							
6		Restricted for MMC-6	Reserved		GROUP NUMBER				
7		(MSB)							
8		VERIFICATION LENGTH							
9		(LSB)							
		CONTROL							

The OPERATION CODE field is defined in SPC-4 and shall be set to the value shown in table 98 for the VERIFY (10) command.

See the READ (10) command (see 5.11) for the definition of the DPO bit. See the PRE-FETCH (10) command (see 5.8) for the definition of the LOGICAL BLOCK ADDRESS field. See the PRE-FETCH (10) command (see 5.8) and 4.23 for the definition of the GROUP NUMBER field.

If the byte check (BYTCHK) field is set to 00b, then:

- a) no Data-Out Buffer transfer shall occur;
- b) for any mapped LBA specified by the command, the device server shall check the protection information from the verify operation based on the VRPROTECT field as defined in table 99; and
- c) for any unmapped LBA specified by the command, the verify operation shall complete without error.

If:

- a) the BYTCHK field is set to 01b or 11b;
- b) the VBULS bit is set to zero in the Block Device Characteristics VPD page (see 6.6.2); and

- c) any LBA specified by the command is unmapped (i.e., deallocated or anchored),

then the device server shall terminate the command with CHECK CONDITION status with the sense key set to MISCOMPARE and the additional sense code set to MISCOMPARE VERIFY OF UNMAPPED LBA.

If:

- a) the BYTCHK field is set to 01b or 11b; and
- b) either:
 - A) the VBULS bit is set to one in the Block Device Characteristics VPD page; or
 - B) all LBAs specified by the command are mapped,

then:

- a) if the BYTCHK field is set to 01b, then the Data-Out Buffer transfer shall include the number of logical blocks specified by the VERIFICATION LENGTH field;
 - b) if the BYTCHK field is set to 11b, then:
 - A) the Data-Out Buffer transfer shall include one logical block; and
 - B) the device server shall:
 - 1) duplicate the single logical block, as described in the WRITE SAME command (see 5.43), the number of times required to satisfy the VERIFICATION LENGTH field; and
 - 2) place the duplicated data in the Data-Out Buffer;
 - c) the device server shall check the protection information transferred from the Data-Out Buffer based on the VRPROTECT field as defined in table 101;
 - d) for any mapped LBA specified by the command, the device server shall perform the verify operation and check the protection information from the verify operation based on the VRPROTECT field as defined in table 100;
- and
- e) the device server shall perform:
 - A) a compare operation of:
 - a) user data from the verify operations; and
 - b) user data from the Data-Out Buffer;
 - and
 - B) a compare operation based on the VRPROTECT field as defined in table 102 of:
 - a) protection information from the verify operations; and
 - b) protection information from the Data-Out Buffer.

The order of the user data and protection information checks and compare operations is vendor specific.

If a compare operation indicates a miscompare, then the device server shall terminate the command with CHECK CONDITION status with the sense key set to MISCOMPARE and the additional sense code set to the appropriate value for the condition.

The VERIFICATION LENGTH field specifies the number of contiguous logical blocks that shall be verified, starting with the logical block referenced by the LBA specified by the LOGICAL BLOCK ADDRESS field. A VERIFICATION LENGTH field set to zero specifies that no logical blocks shall be transferred or verified. This condition shall not be considered an error. If the specified LBA and the specified verification length exceed the capacity of the medium (see 4.5), then the device server shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to LOGICAL BLOCK ADDRESS OUT OF RANGE. The VERIFICATION LENGTH field is constrained by the MAXIMUM TRANSFER LENGTH field in the Block Limits VPD page (see 6.6.3).

If the BYTCHK field is set to 01b, then the VERIFICATION LENGTH field also specifies the number of logical blocks that the device server shall transfer from the Data-Out Buffer.

The CONTROL byte is defined in SAM-5.

If the BYTCHK field is set to 00b, then table 99 defines the checks that the device server shall perform on the protection information from the verify operations based on the VRPROTECT field. All footnotes for table 99 are at the end of the table.

Table 99 – VRPROTECT field with the BYTCHK field set to 00b – checking protection information from the verify operations (part 1 of 3)

Code	Logical unit formatted with protection information	Field in protection information ^g	Extended INQUIRY Data VPD page bit value ^f	If check fails ^{d e} , additional sense code
000b	Yes ⁱ	LOGICAL BLOCK GUARD	GRD_CHK = 1	LOGICAL BLOCK GUARD CHECK FAILED
			GRD_CHK = 0	No check performed
		LOGICAL BLOCK APPLICATION TAG	APP_CHK = 1 ^c	LOGICAL BLOCK APPLICATION TAG CHECK FAILED
			APP_CHK = 0	No check performed
		LOGICAL BLOCK REFERENCE TAG	REF_CHK = 1 ^h	LOGICAL BLOCK REFERENCE TAG CHECK FAILED
			REF_CHK = 0	No check performed
No	No protection information on the medium to check.			
001b 101b ^b	Yes	LOGICAL BLOCK GUARD	GRD_CHK = 1	LOGICAL BLOCK GUARD CHECK FAILED
			GRD_CHK = 0	No check performed
		LOGICAL BLOCK APPLICATION TAG	APP_CHK = 1 ^c	LOGICAL BLOCK APPLICATION TAG CHECK FAILED
			APP_CHK = 0	No check performed
		LOGICAL BLOCK REFERENCE TAG	REF_CHK = 1 ^h	LOGICAL BLOCK REFERENCE TAG CHECK FAILED
			REF_CHK = 0	No check performed
No	Error condition ^a			
010b ^b	Yes	LOGICAL BLOCK GUARD	No check performed	
		LOGICAL BLOCK APPLICATION TAG	APP_CHK = 1 ^c	LOGICAL BLOCK APPLICATION TAG CHECK FAILED
			APP_CHK = 0	No check performed
		LOGICAL BLOCK REFERENCE TAG	REF_CHK = 1 ^h	LOGICAL BLOCK REFERENCE TAG CHECK FAILED
			REF_CHK = 0	No check performed
		No	Error condition ^a	

Table 99 – VRPROTECT field with the BYTCHK field set to 00b – checking protection information from the verify operations (part 2 of 3)

Code	Logical unit formatted with protection information	Field in protection information ^g	Extended INQUIRY Data VPD page bit value ^f	If check fails ^{d e} , additional sense code
011b ^b	Yes	LOGICAL BLOCK GUARD	No check performed	
		LOGICAL BLOCK APPLICATION TAG	No check performed	
		LOGICAL BLOCK REFERENCE TAG	No check performed	
	No	Error condition ^a		
100b ^b	Yes	LOGICAL BLOCK GUARD	GRD_CHK = 1	LOGICAL BLOCK GUARD CHECK FAILED
			GRD_CHK = 0	No check performed
		LOGICAL BLOCK APPLICATION TAG	No check performed	
	LOGICAL BLOCK REFERENCE TAG	No check performed		
No	Error condition ^a			

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC 14776-323:2017

Table 99 – VRPROTECT field with the BYTCHK field set to 00b – checking protection information from the verify operations (part 3 of 3)

Code	Logical unit formatted with protection information	Field in protection information ^g	Extended INQUIRY Data VPD page bit value ^f	If check fails ^{d e} , additional sense code
110b to 111b	Reserved			
<p>^a If the logical unit supports protection information (see 4.22) and has not been formatted with protection information, then the device server shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.</p> <p>^b If the logical unit does not support protection information, then the device server should terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.</p> <p>^c If the device server has knowledge of the contents of the LOGICAL BLOCK APPLICATION TAG field, then the device server shall check each logical block application tag. If the ATO bit in the Control mode page (see SPC-4) is set to one, then this knowledge is acquired from:</p> <ul style="list-style-type: none"> a) the EXPECTED LOGICAL BLOCK APPLICATION TAG field and the LOGICAL BLOCK APPLICATION TAG MASK field in the CDB, if the command is a VERIFY (32) command (see 5.32); b) the Application Tag mode page (see 6.5.3), if the command is not a VERIFY (32) command and the ATMPE bit in the Control mode page (see SPC-4) is set to one; or c) a method not defined by this standard, if the command is not a VERIFY (32) command and the ATMPE bit is set to zero. <p>^d If the device server terminates the command with CHECK CONDITION status, then the device server shall set the sense key to ABORTED COMMAND.</p> <p>^e If multiple errors occur while the device server is processing the command, then the selection by the device server of which error to report is not defined by this standard.</p> <p>^f See the Extended INQUIRY Data VPD page (see SPC-4) for the definitions of the GRD_CHK bit, the APP_CHK bit, and the REF_CHK bit.</p> <p>^g If the device server detects:</p> <ul style="list-style-type: none"> a) a LOGICAL BLOCK APPLICATION TAG field set to FFFFh and type 1 protection (see 4.22.2.3) or type 2 protection (see 4.22.2.4) is enabled; or b) a LOGICAL BLOCK APPLICATION TAG field set to FFFFh, LOGICAL BLOCK REFERENCE TAG field set to FFFF_FFFFh, and type 3 protection (see 4.22.2.5) is enabled, <p>then the device server shall not check any protection information in the associated protection information interval.</p> <p>^h If type 1 protection is enabled, then the device server shall check each logical block reference tag by comparing it to the lower four bytes of the LBA associated with the logical block. If type 2 protection or type 3 protection is enabled, and the device server has knowledge of the contents of the LOGICAL BLOCK REFERENCE TAG field, then the device server shall check the logical block reference tag. If type 2 protection is enabled, then this knowledge may be acquired through the EXPECTED INITIAL LOGICAL BLOCK REFERENCE TAG field in a VERIFY (32) command (see 5.32). If type 3 protection is enabled, then the method for acquiring this knowledge is not defined by this standard.</p> <p>ⁱ If the DPICZ bit in the Control mode page (see SPC-4) is set to one, then protection information shall not be checked.</p>				

If the BYTCHK field is set to 01b or 11b, then table 100 defines the checks that the device server shall perform on the protection information from the verify operations based on the VRPROTECT field. All footnotes for table 100 are at the end of the table.

Table 100 – VRPROTECT field with the BYTCHK field set to 01b or 11b – checking protection information from the verify operations (part 1 of 2)

Code	Logical unit formatted with protection information	Field in protection information ^g	Extended INQUIRY Data VPD page bit value ^f	If check fails ^{d e} , additional sense code
000b	Yes ⁱ	LOGICAL BLOCK GUARD	GRD_CHK = 1	LOGICAL BLOCK GUARD CHECK FAILED
			GRD_CHK = 0	No check performed
		LOGICAL BLOCK APPLICATION TAG	APP_CHK = 1 ^{c g}	LOGICAL BLOCK APPLICATION TAG CHECK FAILED
			APP_CHK = 0	No check performed
		LOGICAL BLOCK REFERENCE TAG	REF_CHK = 1 ^h	LOGICAL BLOCK REFERENCE TAG CHECK FAILED
			REF_CHK = 0	No check performed
No	No protection information available to check			
001b 010b 011b 100b 101b ^b	Yes	LOGICAL BLOCK GUARD	No check performed	
		LOGICAL BLOCK APPLICATION TAG	No check performed	
		LOGICAL BLOCK REFERENCE TAG	No check performed	
	No	Error condition ^a		
110b to 111b	Reserved			

STANDARDSISO.COM Click to view the full PDF of ISO/IEC 14776-323:2017

Table 100 – VRPROTECT field with the BYTCHK field set to 01b or 11b – checking protection information from the verify operations (part 2 of 2)

Code	Logical unit formatted with protection information	Field in protection information ^g	Extended INQUIRY Data VPD page bit value ^f	If check fails ^{d e} , additional sense code
<p>^a If the logical unit supports protection information (see 4.22) and has not been formatted with protection information, then the device server shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.</p> <p>^b If the logical unit does not support protection information, then the device server should terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.</p> <p>^c If the device server has knowledge of the contents of the LOGICAL BLOCK APPLICATION TAG field, then the device server shall check each logical block application tag. If the ATO bit in the Control mode page (see SPC-4) is set to one, then this knowledge is acquired from:</p> <ul style="list-style-type: none"> a) the EXPECTED LOGICAL BLOCK APPLICATION TAG field and the LOGICAL BLOCK APPLICATION TAG MASK field in the CDB, if the command is a VERIFY (32) command (see 5.32); b) the Application Tag mode page (see 6.5.3), if the command is not a VERIFY (32) command and the ATMPE bit in the Control mode page (see SPC-4) is set to one; or c) a method not defined by this standard, if the command is not a VERIFY (32) command and the ATMPE bit is set to zero. <p>^d If the device server terminates the command with CHECK CONDITION status, then the device server shall set the sense key to ABORTED COMMAND.</p> <p>^e If multiple errors occur while the device server is processing the command, then the selection by the device server of which error to report is not defined by this standard.</p> <p>^f See the Extended INQUIRY Data VPD page (see SPC-4) for the definitions of the GRD_CHK bit, the APP_CHK bit, and the REF_CHK bit.</p> <p>^g If the device server detects:</p> <ul style="list-style-type: none"> a) a LOGICAL BLOCK APPLICATION TAG field set to FFFFh and type 1 protection (see 4.22.2.3) or type 2 protection (see 4.22.2.4) is enabled; or b) a LOGICAL BLOCK APPLICATION TAG field set to FFFFh, LOGICAL BLOCK REFERENCE TAG field set to FFFF_FFFFh, and type 3 protection (see 4.22.2.5) is enabled, <p>then the device server shall not check any protection information in the associated protection information interval.</p> <p>^h If type 1 protection is enabled, then the device server shall check each logical block reference tag by comparing it to the lower four bytes of the LBA associated with the logical block. If type 2 protection or type 3 protection is enabled, and the device server has knowledge of the contents of the LOGICAL BLOCK REFERENCE TAG field, then the device server shall check the logical block reference tag. If type 2 protection is enabled, then this knowledge may be acquired through the EXPECTED INITIAL LOGICAL BLOCK REFERENCE TAG field in a VERIFY (32) command (see 5.32). If type 3 protection is enabled, then the method for acquiring this knowledge is not defined by this standard.</p> <p>ⁱ If the DPICZ bit in the Control mode page (see SPC-4) is set to one, then protection information shall not be checked.</p>				

If the BYTCHK field is set to 01b or 11b, then table 101 defines the checks that the device server shall perform on the protection information transferred from the Data-Out Buffer based on the VRPROTECT field. All footnotes for table 101 are at the end of the table.

Table 101 – VRPROTECT field with the BYTCHK field set to 01b or 11b – checking protection information from the Data-Out Buffer (part 1 of 2)

Code	Logical unit formatted with protection information	Field in protection information	Device server check	If check fails ^{d e} , additional sense code
000b	Yes	No protection information in the Data-Out Buffer to check		
	No	No protection information in the Data-Out Buffer to check		
001b ^b	Yes	LOGICAL BLOCK GUARD	Shall	LOGICAL BLOCK GUARD CHECK FAILED
		LOGICAL BLOCK APPLICATION TAG	May ^c	LOGICAL BLOCK APPLICATION TAG CHECK FAILED
		LOGICAL BLOCK REFERENCE TAG	Shall (except for type 3) ^f	LOGICAL BLOCK REFERENCE TAG CHECK FAILED
	No	Error condition ^a		
010b ^b	Yes	LOGICAL BLOCK GUARD	Shall not	No check performed
		LOGICAL BLOCK APPLICATION TAG	May ^c	LOGICAL BLOCK APPLICATION TAG CHECK FAILED
		LOGICAL BLOCK REFERENCE TAG	May ^f	LOGICAL BLOCK REFERENCE TAG CHECK FAILED
	No	Error condition ^a		
011b ^b	Yes	LOGICAL BLOCK GUARD	Shall not	No check performed
		LOGICAL BLOCK APPLICATION TAG	Shall not	No check performed
		LOGICAL BLOCK REFERENCE TAG	Shall not	No check performed
	No	Error condition ^a		
100b ^b	Yes	LOGICAL BLOCK GUARD	Shall	LOGICAL BLOCK GUARD CHECK FAILED
		LOGICAL BLOCK APPLICATION TAG	Shall not	No check performed
		LOGICAL BLOCK REFERENCE TAG	Shall not	No check performed
	No	Error condition ^a		

Table 101 – VRPROTECT field with the BYTCHK field set to 01b or 11b – checking protection information from the Data-Out Buffer (part 2 of 2)

Code	Logical unit formatted with protection information	Field in protection information	Device server check	If check fails ^{d e} , additional sense code
101b ^b	Yes	LOGICAL BLOCK GUARD	Shall	LOGICAL BLOCK GUARD CHECK FAILED
		LOGICAL BLOCK APPLICATION TAG	May ^c	LOGICAL BLOCK APPLICATION TAG CHECK FAILED
		LOGICAL BLOCK REFERENCE TAG	May ^f	LOGICAL BLOCK REFERENCE TAG CHECK FAILED
	No	Error condition ^a		
110b to 111b	Reserved			

^a If the logical unit supports protection information (see 4.22) and has not been formatted with protection information, then the device server shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

^b If the logical unit does not support protection information, then the device server should terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

^c If the device server has knowledge of the contents of the LOGICAL BLOCK APPLICATION TAG field and the ATO bit is set to one in the Control mode page (see SPC-4), then the device server may check each logical block application tag. If the ATO bit is set to one, then this knowledge is acquired from:

- the EXPECTED LOGICAL BLOCK APPLICATION TAG field and the LOGICAL BLOCK APPLICATION TAG MASK field in the CDB, if the command is a VERIFY (32) command;
- the Application Tag mode page (see 6.5.3), if the command is not a VERIFY (32) command and the ATMPE bit in the Control mode page (see SPC-4) is set to one; or
- a method not defined by this standard, if the command is not a VERIFY (32) command and the ATMPE bit is set to zero.

^d If the device server terminates the command with CHECK CONDITION status, then the device server shall set the sense key to ABORTED COMMAND.

^e If multiple errors occur while the device server is processing the command, then the selection by the device server of which error to report is not defined by this standard.

^f If type 1 protection is enabled, then the device server shall check the logical block reference tag by comparing it to the lower four bytes of the LBA associated with the logical block. If type 2 protection is enabled and the device server has knowledge of the contents of the LOGICAL BLOCK REFERENCE TAG field, then the device server shall check each logical block reference tag. If type 2 protection is enabled, then this knowledge may be acquired through the EXPECTED INITIAL LOGICAL BLOCK REFERENCE TAG field in a VERIFY (32) command (see 5.32). If type 3 protection is enabled, the ATO bit is set to one in the Control mode page (see SPC-4), and the device server has knowledge of the contents of the LOGICAL BLOCK REFERENCE TAG field, then the device server may check each logical block reference tag. If type 3 protection is enabled, then the method for acquiring this knowledge is not defined by this standard.

If the BYCHK field is set to 01b or 11b, then table 102 defines the processing by the device server of the protection information during the compare operation based on the VRPROTECT field. All footnotes for table 102 are at the end of the table.

Table 102 – VRPROTECT field with the BYCHK field set to 01b or 11b – compare operation requirements
(part 1 of 3)

Code	Logical unit formatted with protection information	Field in protection information	Compare operation	If compare fails ^{c d} , additional sense code
000b	Yes	No protection information in the Data-Out Buffer to compare. Only user data is compared within each logical block.		
	No	No protection information from the verify operations or in the Data-Out Buffer to compare. Only user data is compared within each logical block.		
001b ^b	Yes	LOGICAL BLOCK GUARD	Shall	LOGICAL BLOCK GUARD CHECK FAILED
		LOGICAL BLOCK APPLICATION TAG (ATO = 1) ^e	Shall	LOGICAL BLOCK APPLICATION TAG CHECK FAILED
		LOGICAL BLOCK APPLICATION TAG (ATO = 0) ^f	Shall not	No compare performed
		LOGICAL BLOCK REFERENCE TAG (not type 3)	Shall ^g	LOGICAL BLOCK REFERENCE TAG CHECK FAILED
		LOGICAL BLOCK REFERENCE TAG (type 3 and ATO = 1)	Shall ^h	LOGICAL BLOCK REFERENCE TAG CHECK FAILED
		LOGICAL BLOCK REFERENCE TAG (type 3 and ATO = 0)	Shall not	No compare performed
	No	Error condition ^a		

STANDARDSISO.COM. Click to view the full PDF of ISO/IEC 14776-323:2017

Table 102 – VRPROTECT field with the BYTCHK field set to 01b or 11b – compare operation requirements
(part 2 of 3)

Code	Logical unit formatted with protection information	Field in protection information	Compare operation	If compare fails ^{c d} , additional sense code
010b ^b	Yes	LOGICAL BLOCK GUARD	Shall not	No compare performed
		LOGICAL BLOCK APPLICATION TAG (ATO = 1) ^e	Shall	LOGICAL BLOCK APPLICATION TAG CHECK FAILED
		LOGICAL BLOCK APPLICATION TAG (ATO = 0) ^f	Shall not	No compare performed
		LOGICAL BLOCK REFERENCE TAG (not type 3)	Shall ^g	LOGICAL BLOCK REFERENCE TAG CHECK FAILED
		LOGICAL BLOCK REFERENCE TAG (type 3 and ATO = 1)	Shall ^h	LOGICAL BLOCK REFERENCE TAG CHECK FAILED
		LOGICAL BLOCK REFERENCE TAG (type 3 and ATO = 0)	Shall not	No compare performed
	No	Error condition ^a		
011b 100b ^b	Yes	LOGICAL BLOCK GUARD	Shall	LOGICAL BLOCK GUARD CHECK FAILED
		LOGICAL BLOCK APPLICATION TAG (ATO = 1) ^e	Shall	LOGICAL BLOCK APPLICATION TAG CHECK FAILED
		LOGICAL BLOCK APPLICATION TAG (ATO = 0) ^f	Shall not	No compare performed
		LOGICAL BLOCK REFERENCE TAG (not type 3)	Shall ^g	LOGICAL BLOCK REFERENCE TAG CHECK FAILED
		LOGICAL BLOCK REFERENCE TAG (type 3 and ATO = 1)	Shall ^h	LOGICAL BLOCK REFERENCE TAG CHECK FAILED
		LOGICAL BLOCK REFERENCE TAG (type 3 and ATO = 0)	Shall not	No compare performed
	No	Error condition ^a		

Table 102 – VRPROTECT field with the BYTCHK field set to 01b or 11b – compare operation requirements
(part 3 of 3)

Code	Logical unit formatted with protection information	Field in protection information	Compare operation	If compare fails ^{c d} , additional sense code
101b ^b	Yes	LOGICAL BLOCK GUARD	Shall	LOGICAL BLOCK GUARD CHECK FAILED
		LOGICAL BLOCK APPLICATION TAG (ATO = 1) ^e	Shall	LOGICAL BLOCK APPLICATION TAG CHECK FAILED
		LOGICAL BLOCK APPLICATION TAG (ATO = 0) ^f	Shall not	No compare performed
		LOGICAL BLOCK REFERENCE TAG	Shall not	No compare performed
	No	Error condition ^a		
110b to 111b	Reserved			

^a If the logical unit supports protection information (see 4.22) and has not been formatted with protection information, then the device server shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

^b If the logical unit does not support protection information, then the device server should terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

^c If the device server terminates the command with CHECK CONDITION status, then the device server shall set the sense key to MISCOMPARE.

^d If multiple errors occur while the device server is processing the command, then the selection by the device server of which error to report is not defined by this standard.

^e If the ATO bit is set to one in the Control mode page (see SPC-4), then the device server shall not modify the logical block application tag.

^f If the ATO bit is set to zero in the Control mode page (see SPC-4), then the device server may modify any logical block application tag.

^g If the BYTCHK field is set to 11b, then the device server shall compare the value from each LOGICAL BLOCK REFERENCE TAG field received in the single logical block data from the Data-Out Buffer with the corresponding LOGICAL BLOCK REFERENCE TAG field in the first logical block from the verify operations, and the device server shall compare the value of the previous LOGICAL BLOCK REFERENCE TAG field plus one with each of the subsequent LOGICAL BLOCK REFERENCE TAG fields (see 4.22.3).

^h If the BYTCHK field is set to 11b, then the device server shall compare the value from each LOGICAL BLOCK REFERENCE TAG field received in the single logical block data from the Data-Out Buffer with the corresponding LOGICAL BLOCK REFERENCE TAG field in each logical block from the verify operations (see 4.22.3).

5.30 VERIFY (12) command

The VERIFY (12) command (see table 103) requests that the device server perform the actions defined for the VERIFY (10) command (see 5.29).

Migration from the VERIFY (12) command to the VERIFY (16) command is recommended for all implementations.

Table 103 – VERIFY (12) command

Byte	Bit	7	6	5	4	3	2	1	0
0		OPERATION CODE (AFh)							
1		VRPROTECT			DPO	Reserved	BYCHK		Obsolete
2	(MSB)	LOGICAL BLOCK ADDRESS							
...									
5									
6	(MSB)	VERIFICATION LENGTH							
...									
9									
10		Restricted for MMC-6	Reserved		GROUP NUMBER				
11		CONTROL							

The OPERATION CODE field is defined in SPC-4 and shall be set to the value shown in table 103 for the VERIFY (12) command.

See the VERIFY (10) command (see 5.29) for the definitions of the other fields in this command.

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC 14776-323:2017

5.31 VERIFY (16) command

The VERIFY (16) command (see table 104) requests that the device server perform the actions defined for the VERIFY (10) command (see 5.29).

Table 104 – VERIFY (16) command

Byte	Bit	7	6	5	4	3	2	1	0
0		OPERATION CODE (8Fh)							
1		VRPROTECT			DPO	Reserved	BYTCHK		Reserved
2	(MSB)	LOGICAL BLOCK ADDRESS							
...									
9									
10	(MSB)	VERIFICATION LENGTH							
...									
13									
14		Restricted for MMC-6	Reserved		GROUP NUMBER				
15		CONTROL							

The OPERATION CODE field is defined in SPC-4 and shall be set to the value shown in table 104 for the VERIFY (16) command.

See the VERIFY (10) command (see 5.29) for the definitions of the other fields in this command.

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC 14776-323:2017

5.32 VERIFY (32) command

The VERIFY (32) command (see table 105) requests that the device server perform the actions defined for the VERIFY (10) command (see 5.29).

The device server shall process a VERIFY (32) command only if type 2 protection is enabled (see 4.22.2.4).

Table 105 – VERIFY (32) command

Byte	Bit	7	6	5	4	3	2	1	0
0		OPERATION CODE (7Fh)							
1		CONTROL							
2		Reserved							
...									
5									
6		Reserved			GROUP NUMBER				
7		ADDITIONAL CDB LENGTH (18h)							
8	(MSB)	SERVICE ACTION (000Ah)							
9									
10		VRPROTECT		DPO	Reserved	BYTCHK		Reserved	
11		Reserved							
12	(MSB)	LOGICAL BLOCK ADDRESS							
...									
19									
20	(MSB)	EXPECTED INITIAL LOGICAL BLOCK REFERENCE TAG							
...									
23									
24	(MSB)	EXPECTED LOGICAL BLOCK APPLICATION TAG							
25									
26	(MSB)	LOGICAL BLOCK APPLICATION TAG MASK							
27									
28	(MSB)	VERIFICATION LENGTH							
...									
31									

The OPERATION CODE field, the ADDITIONAL CDB LENGTH field, and the SERVICE ACTION field are defined in SPC-4 and shall be set to the values shown in table 105 for the VERIFY (32) command.

See the VERIFY (10) command (see 5.29) for the definitions of the CONTROL byte, the GROUP NUMBER field, the VRPROTECT field, the DPO bit, the BYTCHK field, the LOGICAL BLOCK ADDRESS field, and the VERIFICATION LENGTH field.

If checking of the LOGICAL BLOCK REFERENCE TAG field is enabled (see table 99, table 100, table 101, and table 102 in 5.29), then the EXPECTED INITIAL LOGICAL BLOCK REFERENCE TAG field contains the value of the LOGICAL BLOCK REFERENCE TAG field expected in the protection information of the first logical block accessed by the command instead of a value based on the LBA (see 4.22.3).

If the ATO bit is set to one in the Control mode page (see SPC-4) and checking of the LOGICAL BLOCK APPLICATION TAG field is enabled (see table 99, table 100, table 101, and table 102 in 5.29), then the LOGICAL BLOCK APPLICATION TAG MASK field contains a value that is a bit mask for enabling the checking of the LOGICAL BLOCK APPLICATION TAG field in every instance of the protection information for each logical block accessed by the command. A LOGICAL BLOCK APPLICATION TAG MASK bit set to one enables the checking of the corresponding bit of the EXPECTED LOGICAL BLOCK APPLICATION TAG field with the corresponding bit of the LOGICAL BLOCK APPLICATION TAG field in every instance of the protection information. A LOGICAL BLOCK APPLICATION TAG MASK field bit set to zero disables the checking of the corresponding bit of the EXPECTED LOGICAL BLOCK APPLICATION TAG field with the corresponding bit of the LOGICAL BLOCK APPLICATION TAG field in every instance of protection information.

The LOGICAL BLOCK APPLICATION TAG MASK field and the EXPECTED LOGICAL BLOCK APPLICATION TAG field shall be ignored if:

- a) the ATO bit is set to zero; or
- b) the ATO bit is set to one in the Control mode page (see SPC-4) and checking of the LOGICAL BLOCK APPLICATION TAG field is disabled (see table 99, table 100, table 101, and table 102 in 5.29).

5.33 WRITE (10) command

The WRITE (10) command (see table 106) requests that the device server:

- a) transfer the specified logical block data from the Data-Out Buffer; and
- b) perform write operations to the specified LBAs using the transferred logical blocks.

Migration from the WRITE (10) command to the WRITE (16) command is recommended for all implementations.

Table 106 – WRITE (10) command

Byte	Bit	7	6	5	4	3	2	1	0
0		OPERATION CODE (2Ah)							
1		WRPROTECT			DPO	FUA	Reserved	Obsolete	Obsolete
2	(MSB)	LOGICAL BLOCK ADDRESS							
...									
5		(LSB)							
6		Reserved			GROUP NUMBER				
7	(MSB)	TRANSFER LENGTH							
8									
9		CONTROL							

The OPERATION CODE field is defined in SPC-4 and shall be set to the value shown in table 106 for the WRITE (10) command.

The device server shall check the protection information, if any, transferred from the Data-Out Buffer based on

the WRPROTECT field as described in table 107. All footnotes for table 107 are at the end of the table.

Table 107 – WRPROTECT field (part 1 of 2)

Code	Logical unit formatted with protection information	Field in protection information	Device server check	If check fails ^{d i} , additional sense code
000b	Yes ^{f g h}	No protection information received from application client to check		
	No	No protection information received from application client to check		
001b ^b	Yes ^e	LOGICAL BLOCK GUARD	Shall	LOGICAL BLOCK GUARD CHECK FAILED
		LOGICAL BLOCK APPLICATION TAG	Dependent on RWWP ^c	LOGICAL BLOCK APPLICATION TAG CHECK FAILED
		LOGICAL BLOCK REFERENCE TAG	Shall (except for type 3) ^j	LOGICAL BLOCK REFERENCE TAG CHECK FAILED
	No ^a	No protection information available to check		
010b ^b	Yes ^e	LOGICAL BLOCK GUARD	Shall not	No check performed
		LOGICAL BLOCK APPLICATION TAG	Dependent on RWWP ^c	LOGICAL BLOCK APPLICATION TAG CHECK FAILED
		LOGICAL BLOCK REFERENCE TAG	May ^j	LOGICAL BLOCK REFERENCE TAG CHECK FAILED
	No ^a	No protection information available to check		
011b ^b	Yes ^e	LOGICAL BLOCK GUARD	Shall not	No check performed
		LOGICAL BLOCK APPLICATION TAG	Shall not	No check performed
		LOGICAL BLOCK REFERENCE TAG	Shall not	No check performed
	No ^a	No protection information available to check		
100b ^b	Yes ^e	LOGICAL BLOCK GUARD	Shall	LOGICAL BLOCK GUARD CHECK FAILED
		LOGICAL BLOCK APPLICATION TAG	Shall not	No check performed
		LOGICAL BLOCK REFERENCE TAG	Shall not	No check performed
	No ^a	No protection information available to check		
101b ^b	Yes ^e	LOGICAL BLOCK GUARD	Shall	LOGICAL BLOCK GUARD CHECK FAILED
		LOGICAL BLOCK APPLICATION TAG	Dependent on RWWP ^c	LOGICAL BLOCK APPLICATION TAG CHECK FAILED
		LOGICAL BLOCK REFERENCE TAG	May ^j	LOGICAL BLOCK REFERENCE TAG CHECK FAILED
	No ^a	No protection information available to check		

Table 107 – WRPROTECT field (part 2 of 2)

Code	Logical unit formatted with protection information	Field in protection information	Device server check	If check fails ^{d i} , additional sense code
110b to 111b	Reserved			
<p>^a If a logical unit supports protection information (see 4.22) and has not been formatted with protection information, then the device server shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.</p> <p>^b If the logical unit does not support protection information, then the device server should terminate the requested command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.</p> <p>^c If the device server has knowledge of the contents of the LOGICAL BLOCK APPLICATION TAG field and the ATO bit is set to one in the Control mode page (see SPC-4), then the device server:</p> <p>a) may check each logical block application tag if the RWWP bit is set to zero in the Control mode page (see SPC-4); and</p> <p>b) shall check each logical block application tag if the RWWP bit is set to one in the Control mode page. If the ATO bit in the Control mode page (see SPC-4) is set to one, then this knowledge is acquired from:</p> <p>a) the EXPECTED LOGICAL BLOCK APPLICATION TAG field and the LOGICAL BLOCK APPLICATION TAG MASK field in the CDB, if a WRITE (32) command (see 5.36) is received by the device server;</p> <p>b) the Application Tag mode page (see 6.5.3), if a command other than WRITE (32) is received by the device server and the ATMPE bit in the Control mode page (see SPC-4) is set to one; or</p> <p>c) a method not defined by this standard, if a command other than WRITE (32) is received by the device server, and the ATMPE bit is set to zero.</p> <p>^d If the device server terminates the command with CHECK CONDITION status, then the device server shall set the sense key to ABORTED COMMAND.</p> <p>^e The device server shall preserve the contents of protection information (e.g., write it to the medium or store it in non-volatile memory).</p> <p>^f The device server shall write a generated CRC (see 4.22.4.2) into each LOGICAL BLOCK GUARD field.</p> <p>^g If the RWWP bit in the Control mode page (see SPC-4) is set to one, then the device server shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB. If the RWWP bit is set to zero and:</p> <p>a) type 1 protection is enabled, then the device server shall write the least significant four bytes of each LBA into the LOGICAL BLOCK REFERENCE TAG field of each of the written logical blocks; or</p> <p>b) type 2 protection or type 3 protection is enabled, then the device server shall write a value of FFFF_FFFF into the LOGICAL BLOCK REFERENCE TAG field of each of the written logical blocks.</p> <p>^h If the ATO bit is set to one in the Control mode page (see SPC-4), then the device server shall write FFFFh into each LOGICAL BLOCK APPLICATION TAG field. If the ATO bit is set to zero, then the device server may write any value into each LOGICAL BLOCK APPLICATION TAG field.</p> <p>ⁱ If multiple errors occur while the device server is processing the command, then the selection by the device server of which error to report is not defined by this standard.</p> <p>^j If type 1 protection is enabled, then the device server shall check the logical block reference tag by comparing it to the lower four bytes of the LBA associated with the logical block. If type 2 protection is enabled and the device server has knowledge of the contents of the LOGICAL BLOCK REFERENCE TAG field, then the device server shall check each logical block reference tag. If type 2 protection is enabled, then this knowledge may be acquired through the EXPECTED INITIAL LOGICAL BLOCK REFERENCE TAG field in a WRITE (32) command (see 5.36). If type 3 protection is enabled, the ATO bit is set to one in the Control mode page (see SPC-4), and the device server has knowledge of the contents of the LOGICAL BLOCK REFERENCE TAG field, then the device server may check each logical block reference tag. If type 3 protection is enabled, then the method for acquiring this knowledge is not defined by this standard.</p>				

See the READ (10) command (see 5.11) for the definition of the DPO bit.

A force unit access (FUA) bit set to one specifies that the device server shall write the logical blocks to:

- a) the non-volatile cache, if any; or
- b) the medium.

An FUA bit set to zero specifies that the device server shall write the logical blocks to:

- a) volatile cache, if any;
- b) non-volatile cache, if any; or
- c) the medium.

See the PRE-FETCH (10) command (see 5.8) for the definition of the LOGICAL BLOCK ADDRESS field.

See the PRE-FETCH (10) command and 4.23 for the definition of the GROUP NUMBER field.

The TRANSFER LENGTH field specifies the number of contiguous logical blocks of data that shall be transferred from the Data-Out Buffer and written, starting with the logical block referenced by the LBA specified by the LOGICAL BLOCK ADDRESS field. A TRANSFER LENGTH field set to zero specifies that no logical blocks shall be transferred or written. This condition shall not be considered an error. Any other value specifies the number of logical blocks that shall be transferred and written. If the specified LBA and the specified transfer length exceed the capacity of the medium (see 4.5), then the device server shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to LOGICAL BLOCK ADDRESS OUT OF RANGE. The TRANSFER LENGTH field is constrained by the MAXIMUM TRANSFER LENGTH field in the Block Limits VPD page (see 6.6.3).

The CONTROL byte is defined in SAM-5.

5.34 WRITE (12) command

The WRITE (12) command (see table 108) requests that the device server perform the actions defined for the WRITE (10) command (see 5.33).

Migration from the WRITE (12) command to the WRITE (16) command is recommended for all implementations.

Table 108 – WRITE (12) command

Byte	Bit	7	6	5	4	3	2	1	0
0		OPERATION CODE (AAh)							
1		WRPROTECT			DPO	FUA	Reserved	Obsolete	Obsolete
2	(MSB)	LOGICAL BLOCK ADDRESS							
...									
5									
6	(MSB)	TRANSFER LENGTH							
...									
9									
10		Restricted for MMC-6	Reserved		GROUP NUMBER				
11		CONTROL							

The OPERATION CODE field is defined in SPC-4 and shall be set to the value shown in table 108 for the WRITE (12) command.

See the WRITE (10) command (see 5.33) for the definitions of the other fields in this command.

5.35 WRITE (16) command

The WRITE (16) command (see table 109) requests that the device server perform the actions defined for the WRITE (10) command (see 5.33).

Table 109 – WRITE (16) command

Byte	Bit	7	6	5	4	3	2	1	0
0		OPERATION CODE (8Ah)							
1		WRPROTECT			DPO	FUA	Reserved	Obsolete	Reserved
2	(MSB)	LOGICAL BLOCK ADDRESS							
...									
9		(LSB)							
10	(MSB)	TRANSFER LENGTH							
...									
13		(LSB)							
14		Restricted for MMC-6	Reserved			GROUP NUMBER			
15		CONTROL							

The OPERATION CODE field is defined in SPC-4 and shall be set to the value shown in table 109 for the WRITE (16) command.

See the WRITE (10) command (see 5.33) for the definitions of the other fields in this command.

5.36 WRITE (32) command

The WRITE (32) command (see table 110) requests that the device server perform the actions defined for the WRITE (10) command (see 5.33).

The device server shall process a WRITE (32) command only if type 2 protection is enabled (see 4.22.2.4).

Table 110 – WRITE (32) command

Byte	Bit	7	6	5	4	3	2	1	0
0		OPERATION CODE (7Fh)							
1		CONTROL							
2		Reserved							
...									
5									
6		Reserved			GROUP NUMBER				
7		ADDITIONAL CDB LENGTH (18h)							
8	(MSB)	SERVICE ACTION (000Bh)							
9									
10		WRPROTECT		DPO	FUA	Reserved	Obsolete	Reserved	
11		Reserved							
12	(MSB)	LOGICAL BLOCK ADDRESS							
...									
19									
20	(MSB)	EXPECTED INITIAL LOGICAL BLOCK REFERENCE TAG							
...									
23									
24	(MSB)	EXPECTED LOGICAL BLOCK APPLICATION TAG							
...									
25									
26	(MSB)	LOGICAL BLOCK APPLICATION TAG MASK							
...									
27									
28	(MSB)	TRANSFER LENGTH							
...									
31									

The OPERATION CODE field, the ADDITIONAL CDB LENGTH field, and the SERVICE ACTION field are defined in SPC-4 and shall be set to the values shown in table 110 for the WRITE (32) command.

See the WRITE (10) command (see 5.33) for the definitions of the CONTROL byte, the GROUP NUMBER field, the WRPROTECT field, the DPO bit, the FUA bit, the LOGICAL BLOCK ADDRESS field, and the TRANSFER LENGTH field.

If checking of the LOGICAL BLOCK REFERENCE TAG field is enabled (see table 107), then the EXPECTED INITIAL LOGICAL BLOCK REFERENCE TAG field contains the value of the LOGICAL BLOCK REFERENCE TAG field expected in the protection information of the first logical block accessed by the command instead of a value based on the LBA (see 4.22.3).

If the ATO bit is set to one in the Control mode page (see SPC-4) and checking of the LOGICAL BLOCK APPLICATION TAG field is enabled (see table 107), then the LOGICAL BLOCK APPLICATION TAG MASK field contains a value that is a bit mask for enabling the checking of the LOGICAL BLOCK APPLICATION TAG field in every instance of protection information for each logical block accessed by the command. A LOGICAL BLOCK APPLICATION TAG MASK bit set to one enables the checking of the corresponding bit of the EXPECTED LOGICAL BLOCK APPLICATION TAG field with the corresponding bit of the LOGICAL BLOCK APPLICATION TAG field in every instance of protection information. A LOGICAL BLOCK APPLICATION TAG MASK field bit set to zero disables the checking of the corresponding bit of the EXPECTED LOGICAL BLOCK APPLICATION TAG field with the corresponding bit of the LOGICAL BLOCK APPLICATION TAG field in every instance of protection information.

If the ATO bit is set to:

- a) to zero; or
- b) to one in the Control mode page (see SPC-4) and checking of the LOGICAL BLOCK APPLICATION TAG field is disabled (see table 107),

then the LOGICAL BLOCK APPLICATION TAG MASK field and the EXPECTED LOGICAL BLOCK APPLICATION TAG field shall be ignored

5.37 WRITE AND VERIFY (10) command

The WRITE AND VERIFY (10) command (see table 111) requests that the device server:

- 1) transfer the specified the logical block data for the command from the Data-Out Buffer;
- 2) perform write medium operations to the specified LBAs;
- 3) perform verify operations from the specified LBAs; and
- 4) if specified, perform a compare operation on:
 - A) the logical block data transferred from the Data-Out Buffer; and
 - B) the logical block data from the verify operations.

The device server may process the LBAs in any order but shall perform this sequence in the specified order for a given LBA.

Migration from the WRITE AND VERIFY (10) command to the WRITE AND VERIFY (16) command is recommended for all implementations.

Table 111 – WRITE AND VERIFY (10) command

Byte	Bit	7	6	5	4	3	2	1	0
0		OPERATION CODE (2Eh)							
1		WRPROTECT			DPO	Reserved	BYTCHK		Obsolete
2	(MSB)	LOGICAL BLOCK ADDRESS							
...									
5		(LSB)							
6		Reserved			GROUP NUMBER				
7	(MSB)	TRANSFER LENGTH							
8									
9		CONTROL							

The OPERATION CODE field is defined in SPC-4 and shall be set to the value shown in table 111 for the WRITE AND VERIFY (10) command.

See the PRE-FETCH (10) command (see 5.8) for the definition of the LOGICAL BLOCK ADDRESS field. See the PRE-FETCH (10) command and 4.23 for the definition of the GROUP NUMBER field. See the WRITE (10) command (see 5.33) for the definitions of the CONTROL byte, the TRANSFER LENGTH field and the WRPROTECT field. See the READ (10) command (see 5.11) for the definition of the DPO bit.

See the VERIFY (10) command (see 5.29) for definition of the byte check (BYTCHK) field when set to 00b, 01b, and 10b. For a WRITE AND VERIFY (10) command, a BYTCHK field set to 11b is reserved.

5.38 WRITE AND VERIFY (12) command

The WRITE AND VERIFY (12) command (see table 112) requests that the device server perform the actions defined for the WRITE AND VERIFY (10) command (see 5.37).

Migration from the WRITE AND VERIFY (12) command to the WRITE AND VERIFY (16) command is recommended for all implementations.

Table 112 – WRITE AND VERIFY (12) command

Byte	Bit	7	6	5	4	3	2	1	0
0		OPERATION CODE (AEh)							
1		WRPROTECT			DPO	Reserved		BYTCHK	Obsolete
2	(MSB)	LOGICAL BLOCK ADDRESS							
...									
5		(LSB)							
6	(MSB)	TRANSFER LENGTH							
...									
9		(LSB)							
10		Restricted for MMC-6	Reserved		GROUP NUMBER				
11		CONTROL							

The OPERATION CODE field is defined in SPC-4 and shall be set to the value shown in table 112 for the WRITE AND VERIFY (12) command.

See the WRITE AND VERIFY (10) command (see 5.37) for the definitions of the other fields in this command.

5.39 WRITE AND VERIFY (16) command

The WRITE AND VERIFY (16) command (see table 113) requests that the device server perform the actions defined for the WRITE AND VERIFY (10) command (see 5.37).

Table 113 – WRITE AND VERIFY (16) command

Byte	Bit	7	6	5	4	3	2	1	0
0		OPERATION CODE (8Eh)							
1		WRPROTECT			DPO	Reserved	BYTCHK		Reserved
2		(MSB)							
...		LOGICAL BLOCK ADDRESS							
9		(LSB)							
10		(MSB)							
...		TRANSFER LENGTH							
13		(LSB)							
14		Restricted for MMC-6	Reserved			GROUP NUMBER			
15		CONTROL							

The OPERATION CODE field is defined in SPC-4 and shall be set to the value shown in table 113 for the WRITE AND VERIFY (16) command.

See the WRITE AND VERIFY (10) command (see 5.37) for the definitions of the other fields in this command.

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC 14776-323:2017

5.40 WRITE AND VERIFY (32) command

The WRITE AND VERIFY (32) command (see table 114) requests that the device server perform the actions defined for the WRITE AND VERIFY (10) command (see 5.37).

The device server shall process a WRITE AND VERIFY (32) command only if type 2 protection is enabled (see 4.22.2.4).

Table 114 – WRITE AND VERIFY (32) command

Byte	Bit	7	6	5	4	3	2	1	0
0		OPERATION CODE (7Fh)							
1		CONTROL							
2		Reserved							
...									
5		Reserved							
6									
6		Reserved			GROUP NUMBER				
7		ADDITIONAL CDB LENGTH (18h)							
8	(MSB)	SERVICE ACTION (000Ch)							
9									
10		WRPROTECT		DPO	Reserved	BYTCHK		Reserved	
11		Reserved							
12	(MSB)	LOGICAL BLOCK ADDRESS							
...									
19		(LSB)							
20	(MSB)	EXPECTED INITIAL LOGICAL BLOCK REFERENCE TAG							
...									
23		(LSB)							
24	(MSB)	EXPECTED LOGICAL BLOCK APPLICATION TAG							
...									
25		(LSB)							
26	(MSB)	LOGICAL BLOCK APPLICATION TAG MASK							
...									
27		(LSB)							
28	(MSB)	TRANSFER LENGTH							
...									
31		(LSB)							

The OPERATION CODE field, the ADDITIONAL CDB LENGTH field, and the SERVICE ACTION field are defined in SPC-4 and shall be set to the values shown in table 114 for the WRITE AND VERIFY (32) command.

See the WRITE AND VERIFY (10) command (see 5.37) for the definitions of the CONTROL byte, the GROUP NUMBER field, the WRPROTECT field, the DPO bit, the BYTCHK field, the LOGICAL BLOCK ADDRESS field, and the TRANSFER LENGTH field.

If checking of the LOGICAL BLOCK REFERENCE TAG field is enabled (see table 107), then the EXPECTED INITIAL LOGICAL BLOCK REFERENCE TAG field contains the value of the LOGICAL BLOCK REFERENCE TAG field expected in the protection information of the first logical block accessed by the command instead of a value based on the LBA (see 4.22.3).

If the ATO bit is set to one in the Control mode page (see SPC-4) and checking of the LOGICAL BLOCK APPLICATION TAG field is enabled (see table 107 in 5.33), then the LOGICAL BLOCK APPLICATION TAG MASK field contains a value that is a bit mask for enabling the checking of the LOGICAL BLOCK APPLICATION TAG field in every instance of protection information for each logical block accessed by the command. A LOGICAL BLOCK APPLICATION TAG MASK bit set to one enables the checking of the corresponding bit of the EXPECTED LOGICAL BLOCK APPLICATION TAG field with the corresponding bit of the LOGICAL BLOCK APPLICATION TAG field in every instance of protection information. A LOGICAL BLOCK APPLICATION TAG MASK field bit set to zero disables the checking of the corresponding bit of the EXPECTED LOGICAL BLOCK APPLICATION TAG field with the corresponding bit of the LOGICAL BLOCK APPLICATION TAG field in every instance of protection information.

If the ATO bit is set to one in the Control mode page (see SPC-4) and checking of the LOGICAL BLOCK APPLICATION TAG field is disabled (see table 107), or if the ATO bit is set to zero, then the LOGICAL BLOCK APPLICATION TAG MASK field and the EXPECTED LOGICAL BLOCK APPLICATION TAG field shall be ignored.

5.41 WRITE LONG (10) command

The WRITE LONG (10) command (see table 115) requests that the device server mark a logical block or physical block as containing a pseudo unrecovered error, or transfer data for a single logical block or physical block from the Data-Out Buffer and write it to the medium. The data written shall be the same length and shall be in the same order as the data returned by the READ LONG (10) command (see 5.19). The device server shall write the logical block or physical block to the medium and shall not complete the command with GOOD status until the logical block has been written on the medium without error. If a cache contains the specified logical block or physical block, then the device server shall invalidate that logical block or physical block in the cache.

Migration from the WRITE LONG (10) command to the WRITE LONG (16) command is recommended for all implementations.

Table 115 – WRITE LONG (10) command

Byte	Bit	7	6	5	4	3	2	1	0
0		OPERATION CODE (3Fh)							
1		COR_DIS	WR_UNCOR	PBLOCK	Reserved			Obsolete	
2		(MSB)							
...		LOGICAL BLOCK ADDRESS							
5		(LSB)							
6		Reserved							
7		(MSB)							
8		BYTE TRANSFER LENGTH							
9		(LSB)							
		CONTROL							

The OPERATION CODE field is defined in SPC-4 and shall be set to the value shown in table 115 for the WRITE LONG (10) command.

The correction disabled (COR_DIS) bit, the write uncorrectable error (WR_UNCOR) bit, and the physical block (PBLOCK) bit are defined in table 116. If there is more than one logical block per physical block (i.e., the LOGICAL BLOCKS PER PHYSICAL BLOCK EXPONENT field in the READ CAPACITY (16) parameter data (see 5.16.2) is set to a non-zero value), then the device server shall support the WR_UNCOR bit and the PBLOCK bit.

Table 116 – COR_DIS bit, WR_UNCOR bit, and PBLOCK bit (part 1 of 2)

COR_DIS	WR_UNCOR	PBLOCK	More than one logical block per physical block	Description
0	0	0	yes or no	Write only the specified logical block using the value in the BYTE TRANSFER LENGTH field.
		1	no	Terminate the WRITE LONG command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.
			yes	Write the entire physical block containing the specified logical block using the value in the BYTE TRANSFER LENGTH field.
0	1	0	yes or no	Mark only the specified logical block as containing a pseudo unrecovered error with correction enabled (see 4.18.2) in a manner that causes the device server to perform the maximum error recovery as defined by the Read-Write Error Recovery mode page (see 6.5.8). Ignore the BYTE TRANSFER LENGTH field, and transfer no data.
		1	no	Terminate the WRITE LONG command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.
			yes	Mark the entire physical block containing the specified logical block as containing a pseudo unrecovered error with correction enabled (i.e., mark all of the logical blocks in the same physical block that contains the specified logical block as containing a pseudo unrecovered error with correction enabled) (see 4.18.2). Ignore the BYTE TRANSFER LENGTH field, and transfer no data.

Key:

yes = The LOGICAL BLOCKS PER PHYSICAL BLOCK EXPONENT field is set to a non-zero value in the READ CAPACITY (16) parameter data (see 5.16.2).

no = The LOGICAL BLOCKS PER PHYSICAL BLOCK EXPONENT field is set to zero in the READ CAPACITY (16) parameter data.

Table 116 – COR_DIS bit, WR_UNCOR bit, and PBLOCK bit (part 2 of 2)

COR_DIS	WR_UNCOR	PBLOCK	More than one logical block per physical block	Description
1	0	0	yes or no	Mark only the specified logical block as containing a pseudo unrecovered error with correction disabled (see 4.18.2). Write only the specified logical block using the value in the BYTE TRANSFER LENGTH field.
			no	Terminate the WRITE LONG command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.
		1	yes	Mark the entire physical block containing the specified logical block as containing a pseudo unrecovered error with correction disabled (i.e., mark all of the logical blocks in the same physical block that contains the specified logical block as containing a pseudo unrecovered error with correction disabled) (see 4.18.2). Write the entire physical block containing the specified logical block using the value in the BYTE TRANSFER LENGTH field.
1	1	0	yes or no	Mark only the specified logical block as containing a pseudo unrecovered error with correction disabled (see 4.18.2). Ignore the BYTE TRANSFER LENGTH field, and transfer no data.
			no	Terminate the WRITE LONG command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.
		1	yes	Mark the entire physical block containing the specified logical block as containing a pseudo unrecovered error with correction disabled (i.e., mark all of the logical blocks in the same physical block that contains the specified logical block as containing a pseudo unrecovered error with correction disabled) (see 4.18.2). Ignore the BYTE TRANSFER LENGTH field, and transfer no data.
Key: yes = The LOGICAL BLOCKS PER PHYSICAL BLOCK EXPONENT field is set to a non-zero value in the READ CAPACITY (16) parameter data (see 5.16.2). no = The LOGICAL BLOCKS PER PHYSICAL BLOCK EXPONENT field is set to zero in the READ CAPACITY (16) parameter data.				

The setting of the CD_SUP bit in the Extended INQUIRY Data VPD page (see SPC-4) indicates whether or not the logical unit supports the COR_DIS bit being set to one, and the setting of the WU_SUP bit indicates whether or not the logical unit supports the WR_UNCOR bit being set to one.

The LOGICAL BLOCK ADDRESS field specifies an LBA. If the specified LBA exceeds the capacity of the medium (see 4.5), then the device server shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to LOGICAL BLOCK ADDRESS OUT OF RANGE.

If table 116 defines that the value in the BYTE TRANSFER LENGTH field is used, then the BYTE TRANSFER LENGTH field specifies the number of bytes of data that the device server shall transfer from the Data-Out Buffer and write to the specified logical block or physical block. If the BYTE TRANSFER LENGTH field is not set to zero and does not match the data length that the device server returns for a READ LONG command (see 5.19 and see 5.20), then the device server shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB. In the sense data (see 4.18 and SPC-4), the ILI bit shall be set to one, and the difference (i.e., residue) of the requested length minus the actual length in bytes shall be reported in the INFORMATION field (see SPC-4). Negative values shall be indicated by two's complement notation. If the BYTE TRANSFER LENGTH field is set to zero, then no bytes shall be written. This condition shall not be considered an error.

The CONTROL byte is defined in SAM-5.

5.42 WRITE LONG (16) command

The WRITE LONG (16) command (see table 117) requests that the device server mark a logical block or physical block as containing an error, or transfer data for a single logical block or physical block from the Data-Out Buffer and write it to the medium. The data written shall be the same length and shall be in the same order as the data returned by the READ LONG (16) command (see 5.20). The device server shall write the logical block or physical block to the medium and shall not complete the command with GOOD status until the logical block has been written on the medium without error. If a cache contains the specified logical block or physical block, then the device server shall invalidate that logical block or physical block in the cache.

This command uses the SERVICE ACTION OUT (16) CDB format (see clause A.2).

Table 117 – WRITE LONG (16) command

Byte	Bit	7	6	5	4	3	2	1	0
0		OPERATION CODE (9Fh)							
1		COR_DIS	WR_UNCOR	PBLOCK	SERVICE ACTION (11h)				
2	(MSB)	LOGICAL BLOCK ADDRESS							
...									
9									
10		Reserved							
11		Reserved							
12	(MSB)	BYTE TRANSFER LENGTH							
13									
14		Reserved							
15		CONTROL							

The OPERATION CODE field and SERVICE ACTION field are defined in SPC-4 and shall be set to the values shown in table 117 for the WRITE LONG (16) command.

See the WRITE LONG (10) command (see 5.41) for the definitions of the fields in this command.

5.43 WRITE SAME (10) command

The WRITE SAME (10) command (see table 118) requests that the device server transfer a single logical block from the Data-Out Buffer and for each LBA in the specified range of LBAs:

- a) perform a write operation using the contents of that logical block; or
- b) perform an unmap operation.

The device server writes (i.e., subsequent read operations behave as if the device server wrote the single block of user data received from the Data-Out Buffer to each logical block without modification (see 4.7.3.4.3)).

If the medium is formatted with protection information and the WRPROTECT field is set to 000b, then the device server shall write the LOGICAL BLOCK GUARD field, APPLICATION TAG field, and LOGICAL BLOCK REFERENCE TAG field (see 4.22) for each logical block as required in the WRPROTECT field equals 000b row of table 110.

If:

- a) the medium is formatted with protection information;
- b) the WRPROTECT field is not set to 000b or a reserved value (see table 107); and
- c) the protection information from the Data-Out Buffer is set to FFFF_FFFF_FFFF_FFFFh,

then the device server shall write FFFF_FFFF_FFFF_FFFFh to the protection information for each logical block.

If:

- a) the medium is formatted with type 1 or type 2 protection information;
- b) the WRPROTECT field is not set to 000b or a reserved value (see table 107); and
- c) the protection information from the Data-Out Buffer is not set to FFFF_FFFF_FFFF_FFFFh,

then:

- a) the device server shall write the value from the LOGICAL BLOCK REFERENCE TAG field (see 4.22) received in the logical block from the Data-Out Buffer into the corresponding LOGICAL BLOCK REFERENCE TAG field of the first logical block written. The device server shall write the value of the previous LOGICAL BLOCK REFERENCE TAG field plus one into each of the subsequent LOGICAL BLOCK REFERENCE TAG fields;
- b) if the ATO bit is set to one in the Control mode page (see SPC-4) and the and the ATMPE bit is set to zero in the Control mode page, then the device server shall write the logical block application tag received in the logical block from the Data-Out Buffer into the corresponding LOGICAL BLOCK APPLICATION TAG field (see 4.22) of each logical block;
- c) if the ATO bit is set to one in the Control mode page and the and the ATMPE bit is set to zero in the Control mode page, then the device server shall write the value defined in the Application Tag mode page (see 6.5.3) into the corresponding LOGICAL BLOCK APPLICATION TAG field of each logical block;
- d) if the ATO bit is set to zero in the Control mode page, then the device server may write any value into the LOGICAL BLOCK APPLICATION TAG field of each logical block; and
- e) the device server shall write the value from the LOGICAL BLOCK GUARD field (see 4.22) received in the logical block from the Data-Out Buffer into the corresponding LOGICAL BLOCK GUARD field of each logical block.

If:

- a) the medium is formatted with type 3 protection information;
- b) the WRPROTECT field is not set to 000b or a reserved value (see table 107); and
- c) the protection information from the Data-Out Buffer is not set to FFFF_FFFF_FFFF_FFFFh,

then:

- a) if the ATO bit is set to one in the Control mode page (see SPC-4), then the device server shall write the value from the LOGICAL BLOCK REFERENCE TAG field and the LOGICAL BLOCK APPLICATION TAG field received in the logical block from the Data-Out Buffer into the corresponding LOGICAL BLOCK REFERENCE TAG field of each logical block;
- b) if the ATO bit is set to zero in the Control mode page, then the device server may write any value into the LOGICAL BLOCK REFERENCE TAG field of each logical block; and
- c) the device server shall write the value from the LOGICAL BLOCK GUARD field and the LOGICAL BLOCK APPLICATION TAG field received in the logical block from the Data-Out Buffer into the corresponding LOGICAL BLOCK GUARD field of each logical block.

Migration from the WRITE SAME (10) command to the WRITE SAME (16) command is recommended for all implementations.

Table 118 – WRITE SAME (10) command

Byte	Bit	7	6	5	4	3	2	1	0
0		OPERATION CODE (41h)							
1		WRPROTECT			ANCHOR	UNMAP	Obsolete		
2	(MSB)	LOGICAL BLOCK ADDRESS							
...									
5									
6		Reserved			GROUP NUMBER				
7	(MSB)	NUMBER OF LOGICAL BLOCKS							
8									
9		CONTROL							

The OPERATION CODE field is defined in SPC-4 and shall be set to the value shown in table 118 for the WRITE SAME (10) command.

See the WRITE (10) command (see 5.33) for the definition of the WRPROTECT field.

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC 14776-323:2017

If the logical unit supports logical block provisioning management (see 4.7.3), then the ANCHOR bit in the CDB, the UNMAP bit in the CDB, and the ANC_SUP bit in the Logical Block Provisioning VPD page (see 6.6.4) determine how the device server processes the command as described in table 119.

Table 119 – UNMAP bit, ANCHOR bit, and ANC_SUP bit relationships

UNMAP bit ^a	ANCHOR bit	ANC_SUP bit ^b	Action
0	0	0 or 1	Write ^c
	1	0 or 1	Error ^d
1	0	0 or 1	Deallocate request (see 4.7.3.4.3)
		0	Error ^d
	1	1	Anchor request (see 4.7.3.4.3)

^a The device server in a logical unit that supports logical block provisioning management (see 4.7.3) may implement the UNMAP bit.

^b See the Logical Block Provisioning VPD page (see 6.6.4).

^c The device server shall perform the specified write operation to each LBA specified by the command.

^d The device server shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The device server shall ignore the UNMAP bit and the ANCHOR bit, or the device server shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB if:

- the logical unit is fully provisioned (i.e., the LBPME bit is set to zero in the READ CAPACITY (16) parameter data (see 5.16.2)); and
- the UNMAP bit is set to one or the ANCHOR bit is set to one.

See the PRE-FETCH (10) command (see 5.8) for the definition of the LOGICAL BLOCK ADDRESS field.

See the PRE-FETCH (10) command (see 5.8) and 4.23 for the definition of the GROUP NUMBER field.

If the WSNZ bit is set to zero in the Block Limits VPD page (see 6.6.3), then a NUMBER OF LOGICAL BLOCKS field set to a non-zero value specifies the number of contiguous logical blocks that are requested be unmapped or written, starting with the logical block referenced by the LBA specified by the LOGICAL BLOCK ADDRESS field.

If the WSNZ bit is set to zero, then a NUMBER OF LOGICAL BLOCKS field set to zero specifies that the number of contiguous logical blocks that are requested to be unmapped or written includes all of the logical blocks starting with the LBA specified in the LOGICAL BLOCK ADDRESS field to the last logical block on the medium. If the WSNZ bit is set to one and the NUMBER OF LOGICAL BLOCKS field is set to zero, then the device server shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

If the number of logical blocks specified to be unmapped or written exceeds the value indicated in the MAXIMUM WRITE SAME LENGTH field in the Block Limits VPD page (see 6.6.3), then the device server shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

See the WRITE (10) command (see 5.33) for the definition of the CONTROL byte.

5.44 WRITE SAME (16) command

The WRITE SAME (16) command (see table 120) requests that the device server perform the actions defined for the WRITE SAME (10) command (see 5.43).

Table 120 – WRITE SAME (16) command

Byte	Bit	7	6	5	4	3	2	1	0
0		OPERATION CODE (93h)							
1		WRPROTECT			ANCHOR	UNMAP	Obsolete		NDOB
2	(MSB)	LOGICAL BLOCK ADDRESS							
...									
9		(LSB)							
10	(MSB)	NUMBER OF LOGICAL BLOCKS							
...									
13		(LSB)							
14		Reserved			GROUP NUMBER				
15		CONTROL							

The OPERATION CODE field is defined in SPC-4 and shall be set to the value shown in table 120 for the WRITE SAME (16) command.

The no Data-Out Buffer (NDOB) bit specifies if data is transferred from the Data-Out Buffer. Table 121 defines the interactions between the NDOB bit and the UNMAP bit.

Table 121 – NDOB bit and UNMAP bit interactions

NDOB	UNMAP	Description
0	0 or 1	The device server shall process the command using logical block data from the Data-Out Buffer.
1	0	The device server shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.
	1	The device server shall not transfer data from the Data-Out Buffer and shall process the command as if the Data-Out Buffer contained user data set to all zeroes and protection information, if any, containing: <ul style="list-style-type: none"> a) the LOGICAL BLOCK GUARD field set to FFFFh; b) the LOGICAL BLOCK REFERENCE TAG field set to FFFF_FFFFh; and c) the LOGICAL BLOCK APPLICATION TAG field set to FFFFh.

See the WRITE SAME (10) command (see 5.43) for the definitions of the other fields in this command.

5.45 WRITE SAME (32) command

The WRITE SAME (32) command (see table 122) requests that the device server perform the actions defined for the WRITE SAME (10) command (see 5.43).

The device server shall process a WRITE SAME (32) command only if type 2 protection is enabled (see 4.22.2.4).

Table 122 – WRITE SAME (32) command

Byte	Bit	7	6	5	4	3	2	1	0
0		OPERATION CODE (7Fh)							
1		CONTROL							
2		Reserved							
...									
5									
6		Reserved			GROUP NUMBER				
7		ADDITIONAL CDB LENGTH (18h)							
8	(MSB)	SERVICE ACTION (000Dh)							
9									
10		WRPROTECT			ANCHOR	UNMAP	Obsolete		NDOB
11		Reserved							
12	(MSB)	LOGICAL BLOCK ADDRESS							
...									
19									
20	(MSB)	EXPECTED INITIAL LOGICAL BLOCK REFERENCE TAG							
...									
23									
24	(MSB)	EXPECTED LOGICAL BLOCK APPLICATION TAG							
25									
26	(MSB)	LOGICAL BLOCK APPLICATION TAG MASK							
27									
28	(MSB)	NUMBER OF LOGICAL BLOCKS							
...									
31									

The OPERATION CODE field, the ADDITIONAL CDB LENGTH field, and the SERVICE ACTION field are defined in SPC-4 and shall be set to the values shown in table 122 for the WRITE SAME (32) command.

See the WRITE SAME (10) command (see 5.43) for the definitions of the CONTROL byte, the GROUP NUMBER field, the WRPROTECT field, the LOGICAL BLOCK ADDRESS field, the NUMBER OF LOGICAL BLOCKS field, the UNMAP bit, and the ANCHOR bit.

See the WRITE SAME (16) command (see 5.44) for the definition of the NDOB bit.

If checking of the LOGICAL BLOCK REFERENCE TAG field is enabled (see table 107 in 5.33), then the EXPECTED INITIAL LOGICAL BLOCK REFERENCE TAG field contains the value of the LOGICAL BLOCK REFERENCE TAG field expected in the protection information of the first logical block accessed by the command instead of a value based on the LBA (see 4.22.3).

If the ATO bit is set to one in the Control mode page (see SPC-4) and checking of the LOGICAL BLOCK APPLICATION TAG field is enabled (see table 107 in 5.33), then the LOGICAL BLOCK APPLICATION TAG MASK field contains a value that is a bit mask for enabling the checking of the LOGICAL BLOCK APPLICATION TAG field in every instance of protection information for each logical block accessed by the command. A LOGICAL BLOCK APPLICATION TAG MASK bit set to one enables the checking of the corresponding bit of the EXPECTED LOGICAL

BLOCK APPLICATION TAG field with the corresponding bit of the LOGICAL BLOCK APPLICATION TAG field in every instance of protection information. A LOGICAL BLOCK APPLICATION TAG MASK field bit set to zero disables the checking of the corresponding bit of the EXPECTED LOGICAL BLOCK APPLICATION TAG field with the corresponding bit of the LOGICAL BLOCK APPLICATION TAG field in every instance of protection information.

If the ATO bit is set to one in the Control mode page (see SPC-4) and checking of the LOGICAL BLOCK APPLICATION TAG field is disabled (see table 107 in 5.33), or if the ATO bit is set to zero, then the LOGICAL BLOCK APPLICATION TAG MASK field and the EXPECTED LOGICAL BLOCK APPLICATION TAG field shall be ignored.

5.46 WRITE USING TOKEN command

5.46.1 WRITE USING TOKEN command overview

The WRITE USING TOKEN command (see table 123) requests that the copy manager (see SPC-4) write logical block data represented by the specified ROD token to the specified LBAs.

Table 123 – WRITE USING TOKEN command

Byte	Bit	7	6	5	4	3	2	1	0	
0		OPERATION CODE (83h)								
1		Reserved			SERVICE ACTION (11h)					
2		Reserved								
...										
5		LIST IDENTIFIER								
6	(MSB)									
...										
9		(LSB)								
10	(MSB)	PARAMETER LIST LENGTH								
...										
13										(LSB)
14		Reserved			GROUP NUMBER					
15		CONTROL								

The OPERATION CODE field and the SERVICE ACTION field are defined in SPC-4 and shall be set to the values shown in table 123 for the WRITE USING TOKEN command.

The LIST IDENTIFIER field is defined in SPC-4. The list identifier shall be processed as if the LIST ID USAGE field is set to 00b in the parameter data for an EXTENDED COPY(LID4) command (see SPC-4).

The PARAMETER LIST LENGTH field specifies the length in bytes of the parameter data that is available to be transferred from the Data-Out Buffer. A PARAMETER LIST LENGTH set to zero specifies that no data shall be transferred. This shall not be considered an error.

See the PRE-FETCH (10)FETCH (10) command (see 5.8) and 4.23 for the definition of the GROUP NUMBER field.

The CONTROL byte is defined in SAM-5.

5.46.2 WRITE USING TOKEN parameter list

The parameter list for the WRITE USING TOKEN command is defined in table 124.

Table 124 – WRITE USING TOKEN parameter list

Byte	Bit	7	6	5	4	3	2	1	0
0	(MSB)	WRITE USING TOKEN DATA LENGTH (n - 1)							
1		(LSB)							
2		Reserved						DEL_TKN	IMMED
3		Reserved							
...		Reserved							
7		Reserved							
8	(MSB)	OFFSET INTO ROD							
...		OFFSET INTO ROD							
15		(LSB)							
16		ROD TOKEN							
...		ROD TOKEN							
527		ROD TOKEN							
528		Reserved							
...		Reserved							
533		Reserved							
534		BLOCK DEVICE RANGE DESCRIPTOR LENGTH (n - 535)							
535		BLOCK DEVICE RANGE DESCRIPTOR LENGTH (n - 535)							
Block device range descriptor list (if any)									
536		Block device range descriptor [first] (see 5.7.3)							
...		Block device range descriptor [first] (see 5.7.3)							
551		Block device range descriptor [first] (see 5.7.3)							
		⋮							
n - 15		Block device range descriptor [last] (see 5.7.3)							
...		Block device range descriptor [last] (see 5.7.3)							
n		Block device range descriptor [last] (see 5.7.3)							

The WRITE USING TOKEN DATA LENGTH field specifies the length in bytes of the data that is available to be transferred from the Data-Out Buffer. The write using token data length does not include the number of bytes in the WRITE USING TOKEN DATA LENGTH field.

If the WRITE USING TOKEN DATA LENGTH field is less than 0226h (i.e., 550), then the copy manager shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST.

If the contents of the WRITE USING TOKEN DATA LENGTH field is not equal to the contents of the BLOCK DEVICE RANGE DESCRIPTOR LENGTH field plus 534 then the device server should terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST.

The immediate (IMMED) bit specifies when the copy manager shall return status for the WRITE USING TOKEN command. If the IMMED bit is set to zero, then the copy manager shall process the WRITE USING TOKEN

command until all specified operations are complete or an error is detected. If the IMMED bit is set to one, then the copy manager:

- 1) shall validate the CDB (i.e., detect and report all errors in the CDB);
- 2) shall transfer all the parameter data to the copy manager;
- 3) may validate the parameter data;
- 4) shall complete the WRITE USING TOKEN command with GOOD status; and
- 5) shall complete processing of all specified operations as a background operation (see SPC-4).

If the operations specified by a WRITE USING TOKEN command are processed as a background operation (i.e., the IMMED bit is set to one) (see SPC-4), then the copy manager shall not generate deferred errors (see SAM-5) to report the errors encountered, if any, during this processing. The copy manager shall make error information available to an application client using a RECEIVE ROD TOKEN INFORMATION command (see 5.22).

A delete token (DEL_TKN) bit set to one specifies that the ROD token specified in the ROD TOKEN field should be deleted when processing of the WRITE USING TOKEN command is complete. A DEL_TKN bit set to zero specifies that the ROD token lifetime for the ROD token specified in the ROD TOKEN field shall be as described in SPC-4.

The OFFSET INTO ROD field specifies the offset into the data represented by the ROD token from the first byte represented by the ROD token to the first byte to be transferred. The offset is specified in number of blocks based on the logical block length of the logical unit to which the WRITE USING TOKEN command is to write data. The copy manager that processes the WRITE USING TOKEN command shall compute the byte offset into the ROD by multiplying the contents of the OFFSET INTO ROD field by the logical block length of the logical unit to which the WRITE USING TOKEN command is to write data.

EXAMPLE - To calculate an offset, a ROD token is created from LBAs 15 to 20 followed by LBAs 40 to 100 by a copy manager associated with a logical unit with a logical block length of 512 bytes per logical block. That ROD token is specified in a WRITE USING TOKEN command that transfers one logical block to a logical unit with a logical block length of 4 096 bytes per logical block. The subsequent RECEIVE ROD TOKEN INFORMATION command indicates the successful transfer of 4 096 bytes by setting the TRANSFER COUNT field to one. To create a WRITE USING TOKEN command that transfers bytes from the ROD token starting at the point where the previous WRITE USING TOKEN command stopped, the OFFSET INTO ROD field is set to one (i.e., the contents of the TRANSFER COUNT field) plus the value in the OFFSET INTO ROD field in the previous WRITE USING TOKEN command (i.e., zero). The copy manager multiplies one (i.e., the value in the OFFSET INTO ROD field) by 4 096 (i.e., the logical block length for the logical unit to which the data is being written) and the result is 4 096. As a result, the ROD token logical block that is the start of the transfer is LBA 8 (i.e., LBA 42 from the logical unit whose logical block length is 512 bytes per logical block that was used to create the ROD token).

If the computed byte offset into the ROD is greater than or equal to the number of bytes represented by the ROD token, then the copy manager shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST.

If the ROD TOKEN LENGTH field (see SPC-4) in the ROD TOKEN field is not set to 01F8h, then the copy manager shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense key set to INVALID TOKEN OPERATION, INVALID TOKEN LENGTH.

The ROD TOKEN field specifies the ROD token that represents the data from which logical block data is written. The ROD token is defined as follows:

- a) a ROD token returned by a RECEIVE ROD TOKEN INFORMATION command; or
- b) a block device zero ROD token (see 4.30.4).

If the ROD token does not match any known to the copy manager, then the copy manager shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID TOKEN OPERATION, TOKEN UNKNOWN.

The BLOCK DEVICE RANGE DESCRIPTOR LENGTH field specifies the length in bytes of the block device range descriptor list. The block device range descriptor list length should be a multiple of 16. If the block device range descriptor list length is not a multiple of 16, then the last block device range (see 5.7.3) descriptor is incomplete and shall be ignored. If the block device range descriptor list length is less than 16, then the copy

manager shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST.

If the number of complete block device range descriptors is larger than the MAXIMUM RANGE DESCRIPTORS field in the Block Device ROD Token Limits descriptor (see 6.6.6.3), then the copy manager shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to TOO MANY SEGMENT DESCRIPTORS.

If the same LBA is included in more than one block device range descriptor, then the copy manager shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST.

If the number of bytes of user data represented by the sum of the contents of the NUMBER OF LOGICAL BLOCKS fields in all of the complete block device range descriptors is larger than:

- a) the MAXIMUM BYTES IN BLOCK ROD field in the block ROD device type specific features descriptor in the ROD token features third-party copy descriptor in the Third-party Copy VPD page (see SPC-4) and that field is set to a nonzero value; or
- b) the MAXIMUM TOKEN TRANSFER SIZE field in the Block Device ROD Token Limits descriptor (see 6.6.6.3) and that field is set to a nonzero value,

then the copy manager shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST.

If the number of bytes of user data represented by the sum of the contents of the NUMBER OF LOGICAL BLOCKS fields in all of the complete block device range descriptors is larger than the number of bytes in the data represented by the ROD token minus the computed byte offset to the ROD (i.e., the total requested length of the transfer exceeds the length of the data available in the data represented by the ROD token), then the copy manager shall:

- a) transfer as many whole logical blocks as possible; and
- b) if any portion of a logical block that is written by the copy manager corresponds to offsets into the ROD at or beyond the length of the data represented by the ROD token, then write that portion of the logical block with user data with all bits set to zero.

The copy manager may perform this check during the processing of each block device range descriptor.

5.47 XDWRITEREAD (10) command

The XDWRITEREAD (10) command (see table 125) requests that the device server perform the following as an uninterrupted sequence of actions (see 4.27):

- 1) perform read operations from the specified LBAs;
- 2) transfer the specified number of logical blocks from the Data-Out Buffer;
- 3) perform an XOR operation on:
 - A) the user data contained in the logical blocks from the read operations; and
 - B) the user data contained in the logical blocks transferred from the Data-Out Buffer;
- 4) store the results of the XOR operation (i.e., the XOR data) in a buffer;
- 5) if the DISABLE WRITE bit is set to zero, then perform write operations to the specified LBAs using the logical block data transferred from the Data-Out Buffer; and
- 6) transfer the results of the XOR operation to the Data-In Buffer.

The device server may process the LBAs in any order but shall perform this sequence in the specified order for a given LBA.

Logical block data for this command may include protection information, based on the WRPROTECT field, the XORPINFO bit, and the medium format. This command is only available on transport protocols supporting bidirectional commands.

Table 125 – XDWRITEREAD (10) command

Byte	Bit	7	6	5	4	3	2	1	0
0		OPERATION CODE (53h)							
1		WRPROTECT			DPO	FUA	DISABLE WRITE	Obsolete	XORPINFO
2	(MSB)	LOGICAL BLOCK ADDRESS							
...									
5									
6		Reserved			GROUP NUMBER				
7	(MSB)	TRANSFER LENGTH							
8									
9		CONTROL							

The OPERATION CODE field is defined in SPC-4 and shall be set to the value shown in table 125 for the XDWRITEREAD (10) command.

See the WRITE (10) command (see 5.33) for the definitions of the WRPROTECT field and the FUA bit.

See the READ (10) command (see 5.11) for the definition of the DPO bit.

A DISABLE WRITE bit set to zero specifies that the device server shall perform a write operation using data transferred from the Data-Out Buffer after the XOR operation is complete. A DISABLE WRITE bit set to one specifies that the device server shall not perform a write operation.

If the XOR protection information (XORPINFO) bit is set to zero, then the device server shall not check or transmit protection information.

If the XORPINFO bit is set to one, the device server supports protection information, and the medium has been formatted with protection information, then the device server shall transmit protection information but shall not check any of the protection information fields.

If the XORPINFO bit is set to one and the device server:

- a) supports protection information and the medium has not been formatted with protection information or
- b) does not support protection information,

then the device server shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

See the PRE-FETCH (10) command (see 5.8) for the definition of the LOGICAL BLOCK ADDRESS field.

See the PRE-FETCH (10) command (see 5.8) and 4.23 for the definition of the GROUP NUMBER field.

The TRANSFER LENGTH field specifies the number of contiguous logical blocks of data that the device server shall read, transfer from the Data-Out Buffer, and XOR into a buffer, starting with the logical block referenced by the LBA specified by the LOGICAL BLOCK ADDRESS field. If the specified LBA and the specified transfer length exceed the capacity of the medium (see 4.5), then the device server shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to LOGICAL BLOCK ADDRESS OUT OF RANGE. The TRANSFER LENGTH field is constrained by the MAXIMUM TRANSFER LENGTH field in the Block Limits VPD page (see 6.6.4).

The CONTROL byte is defined in SAM-5.

5.48 XDWRITEREAD (32) command

The XDWRITEREAD (32) command (see table 126) requests that the device server perform the actions defined for the XDWRITEREAD (10) command (see 5.47).

Table 126 – XDWRITEREAD (32) command

Byte	Bit	7	6	5	4	3	2	1	0
0		OPERATION CODE (7Fh)							
1		CONTROL							
2		Reserved							
...									
5									
6		Reserved			GROUP NUMBER				
7		ADDITIONAL CDB LENGTH (18h)							
8	(MSB)	SERVICE ACTION (0007h)							
9									
10		WRPROTECT			DPO	FUA	DISABLE WRITE	Obsolete	XORPINFO
11		Reserved							
12	(MSB)	LOGICAL BLOCK ADDRESS							
...									
19									
20		Reserved							
...									
27									
28	(MSB)	TRANSFER LENGTH							
...									
31									

The OPERATION CODE field, the ADDITIONAL CDB LENGTH field, and the SERVICE ACTION field are defined in SPC-4 and shall be set to the values shown in table 126 for the XDWRITEREAD (32) command.

See the XDWRITEREAD (10) command (see 5.47) for the definitions of the other fields in this command.

5.49 XPWRITE (10) command

The XPWRITE (10) command (see table 127) requests that the device server perform the following as an uninterrupted sequence of actions (see 4.27):

- 1) perform read operations from the specified LBAs;
 - 2) transfer the specified number of logical blocks from the Data-Out Buffer;
 - 3) perform an XOR operation on:
 - A) the user data contained in the logical block data from the read operations; and
 - B) the user data contained in the logical block data transferred from the Data-Out Buffer;
- and

- 4) perform write operations to the specified LBAs using the results of the XOR operation (i.e., the XOR data).

The device server may process the LBAs in any order but shall perform this sequence in the specified order for a given LBA.

Logical block data for this command may include protection information, based on the XORPINFO bit and the medium format.

Table 127 – XPWRITE (10) command

Byte	Bit	7	6	5	4	3	2	1	0	
0		OPERATION CODE (51h)								
1		Reserved			DPO	FUA	Reserved	Obsolete	XORPINFO	
2	(MSB)	LOGICAL BLOCK ADDRESS								
...										
5										(LSB)
6		Reserved			GROUP NUMBER					
7	(MSB)	TRANSFER LENGTH								
8										(LSB)
9										CONTROL

The OPERATION CODE field is defined in SPC-4 and shall be set to the value shown in table 127 for the XPWRITE (10) command.

See the READ (10) command (see 5.11) for the definition of the DPO bit. See the WRITE (10) command (see 5.33) for the definition of the FUA bit. See the PRE-FETCH (10) command (see 5.8) for the definition of the LOGICAL BLOCK ADDRESS field. See the PRE-FETCH (10) command (see 5.8) and 4.23 for the definition of the GROUP NUMBER field.

If the XOR protection information (XORPINFO) bit is set to zero, the device server supports protection information, and the medium has been formatted with protection information, then the device server shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

If the XORPINFO bit is set to one, the device server supports protection information, and the medium has been formatted with protection information, then the device server shall XOR the logical block data transferred from the Data-Out Buffer with the logical block data read, and then write the resulting XOR data. The device server shall not check any of the protection information fields.

If the XORPINFO bit is set to one and the device server:

- a) supports protection information and the medium has not been formatted with protection information;
- or
- b) does not support protection information,

then the device server shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.

The TRANSFER LENGTH field specifies the number of contiguous logical blocks that shall be read, XORed with logical blocks transferred from the Data-Out Buffer, and written, starting with the logical block referenced by the LBA specified by the LOGICAL BLOCK ADDRESS field. If the specified LBA and the specified transfer length exceed the capacity of the medium (see 4.5), then the device server shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to LOGICAL BLOCK ADDRESS OUT OF RANGE. The TRANSFER LENGTH field is constrained by the MAXIMUM TRANSFER LENGTH field in the Block Limits VPD page (see 6.6.3).

The CONTROL byte is defined in SAM-5.

5.50 XPWRITE (32) command

The XPWRITE (32) command (see table 128) requests that the device server perform the actions defined for the XPWRITE (10) command (see 5.49).

Table 128 – XPWRITE (32) command

Byte	Bit	7	6	5	4	3	2	1	0	
0	OPERATION CODE (7Fh)									
1	CONTROL									
2	Reserved									
...										
5										
6	Reserved			GROUP NUMBER						
7	ADDITIONAL CDB LENGTH (18h)									
8	(MSB)	SERVICE ACTION (0006h)							(LSB)	
9										
10	Reserved			DPO	FUA	Reserved	Obsolete	XORPINFO		
11	Reserved									
12	(MSB)	LOGICAL BLOCK ADDRESS							(LSB)	
...										
19										
20	Reserved									
...										
27										
28	(MSB)	TRANSFER LENGTH							(LSB)	
...										
31										

The OPERATION CODE field, the ADDITIONAL CDB LENGTH field, and the SERVICE ACTION field are defined in SPC-4 and shall be set to the values shown in table 128 for the XPWRITE (32) command.

See the XPWRITE (10) command (see 5.49) for the definitions of the other fields in this command.

6 Parameters for direct access block devices

6.1 Parameters for direct access block devices introduction

Table 129 shows the parameters for direct access block devices defined in clause 6 and a reference to the subclause where each parameter type is defined.

Table 129 – Parameters for direct access block devices

Parameter type	Reference
Address descriptors	6.2
Diagnostic parameters	6.3
Log parameters	6.4
Mode parameters	6.5
Vital product data (VPD) parameters	6.6
Copy manager parameters	6.7

6.2 Address descriptors

6.2.1 Address descriptor overview

This subclause describes the address descriptors (see table 130) used for:

- the FORMAT UNIT command (see 5.3);
- the READ DEFECT DATA commands (see 5.17 and 5.18);
- the Translate Address Input diagnostic page (see 6.3.4) for the SEND DIAGNOSTIC command (see SPC-4); and
- the Translate Address Output diagnostic page (see 6.3.5) for the RECEIVE DIAGNOSTIC RESULTS command (see SPC-4).

The format type of an address descriptor is:

- specified in the DEFECT LIST FORMAT field in the CDB for the FORMAT UNIT command (see 5.3.1);
- indicated in the DEFECT LIST FORMAT field in the READ DEFECT DATA parameter data (see 5.17.2 and 5.18.2);
- specified in the SUPPLIED FORMAT field and the TRANSLATE FORMAT field for the Translate Address Output diagnostic page; or
- indicated in the SUPPLIED FORMAT field and the TRANSLATE FORMAT field for the Translate Address Input diagnostic page.

Table 130 defines the types of address descriptors.

Table 130 – Address descriptors

Type	Description	Reference
000b	Short block format address descriptor	6.2.2
001b	Extended bytes from index format address descriptor ^a	6.2.3
010b	Extended physical sector format address descriptor ^a	6.2.4
011b	Long block format address descriptor	6.2.5
100b	Bytes from index format address descriptor ^a	6.2.6
101b	Physical sector format address descriptor ^a	6.2.7
110b	Vendor specific	
111b	Reserved	

^a This address descriptor format type is defined for direct access block devices using rotating media (see 4.3.2).

6.2.2 Short block format address descriptor

A format type of 000b specifies the short block format address descriptor (see table 131).

Table 131 – Short block format address descriptor (000b)

Byte	Bit	7	6	5	4	3	2	1	0
0	(MSB)	SHORT BLOCK ADDRESS							
...									
3	(LSB)								

For the FORMAT UNIT parameter list, the SHORT BLOCK ADDRESS field specifies a four-byte LBA. If the physical block containing the logical block referenced by the specified LBA contains additional logical blocks, then the device server may consider the LBAs of those additional logical blocks to also have been specified.

For the READ DEFECT DATA parameter data, the SHORT BLOCK ADDRESS field indicates a vendor specific four-byte value.

For the Translate Address diagnostic pages, the SHORT BLOCK ADDRESS field contains:

- a) a four-byte LBA, if the value is less than or equal to the capacity of the medium; or
- b) a vendor specific four-byte value, if the value is greater than the capacity of the medium.

6.2.3 Extended bytes from index address descriptor

A format type of 001b specifies the extended bytes from index format address descriptor (see table 132). For the FORMAT UNIT parameter list and the READ DEFECT DATA parameter data, this address descriptor contains the location of a defect that:

- a) is the length of one track (see 4.3.2);
- b) is less than the length of a physical block; or
- c) starts from one address descriptor and extends to the next address descriptor.

For the Translate Address diagnostic pages, this address descriptor contains the location of an LBA. For the Translate Address Output diagnostic page (see 6.2.5), if the SUPPLIED FORMAT field is set to 001b and the MADS bit in the ADDRESS TO TRANSLATE field is set to one, then the device server shall terminate the SEND DIAGNOSTIC command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST.

Table 132 – Extended bytes from index format address descriptor (001b)

Byte	Bit	7	6	5	4	3	2	1	0
0	(MSB)	CYLINDER NUMBER							
...									
2									
3		HEAD NUMBER							
4	MADS	Reserved			(MSB)				
...		BYTES FROM INDEX							
7									

The CYLINDER NUMBER field contains the cylinder number (see 4.3.2).

The HEAD NUMBER field contains the head number (see 4.3.2).

A multi-address descriptor start (MADS) bit set to one specifies that this address descriptor defines the beginning of a defect that spans multiple addresses. The defect may be a number of sequential physical blocks on the same cylinder and head (i.e., a track) or may span a number of sequential tracks on the same head. A MADS bit set to zero specifies that:

- a) this address descriptor defines the end of a defect if the previous address descriptor has the MADS bit set to one; or
- b) this address descriptor defines a single track that contains one or more defects (i.e., the BYTES FROM INDEX field contains FFF_FFFFh) or a single defect (i.e., the BYTES FROM INDEX field does not contain FFF_FFFFh).

See 4.13.2 for valid combinations of two address descriptors that describe a defect.

The BYTES FROM INDEX field:

- a) if not set to FFF_FFFFh, contains the number of bytes from the index (e.g., from the start of the track) to the location being described; or
- b) if set to FFF_FFFFh, specifies or indicates that the entire track is being described.

More than one logical block may be described by this address descriptor.

Table 133 defines the order of the fields used for sorting extended bytes from index format address descriptors if the command using the address descriptors specifies sorting.

Table 133 – Sorting order for extended bytes from index format address descriptors

Bit:	(MSB)	...	36	35	...	28	27	...	(LSB)	
	59								0	
CYLINDER NUMBER field				HEAD NUMBER field				BYTES FROM INDEX field		

6.2.4 Extended physical sector format address descriptor

A format type of 010b specifies the extended physical sector format address descriptor (see table 134). For the FORMAT UNIT parameter list and the READ DEFECT DATA parameter data, this address descriptor contains the location of a defect that:

- is the length of one track (see 4.3.2);
- is less than the length of a physical block; or
- starts from one address descriptor and extends to the next address descriptor.

For the Translate Address diagnostic pages, this address descriptor specifies the location of an LBA. For the Translate Address Output diagnostic page (see 6.2.5), if the SUPPLIED FORMAT field is set to 010b and the MADS bit in the ADDRESS TO TRANSLATE field is set to one, then the device server shall terminate the SEND DIAGNOSTIC command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST.

Table 134 – Extended physical sector format address descriptor (010b)

Byte	Bit	7	6	5	4	3	2	1	0
0	(MSB)	CYLINDER NUMBER							
...									
2	(LSB)								
3		HEAD NUMBER							
4	MADS	Reserved			(MSB)				
...		SECTOR NUMBER							
7	(LSB)								

The CYLINDER NUMBER field contains the cylinder number (see 4.3.2).

The HEAD NUMBER field contains the head number (see 4.3.2).

A multi-address descriptor start (MADS) bit set to one specifies that this address descriptor defines the beginning of a defect that spans multiple addresses. The defect may span a number of sequential physical blocks on the same cylinder and head (i.e., a track) or may span a number of sequential tracks on the same head. A MADS bit set to zero specifies that:

- this address descriptor defines the end of a defect if the previous address descriptor has the MADS bit set to one; or
- this address descriptor defines a single track that contains one or more defects (i.e., the SECTOR NUMBER field contains FFF_FFFFh) or a single defect (i.e., the SECTOR NUMBER field does not contain FFF_FFFFh).

See 4.13.2 for valid combinations of two address descriptors that describe a defect.

The SECTOR NUMBER field:

- if not set to FFF_FFFFh, contains the sector number (see 4.3.2) of the location being described; or
- if set to FFF_FFFFh, specifies or indicates that the entire track is being described.

More than one logical block may be described by this address descriptor.

The HEAD NUMBER field contains the head number (see 4.3.2).

The BYTES FROM INDEX field contains the number of bytes from the index (e.g., from the start of the track) to the location being described. A BYTES FROM INDEX field set to FFFF_FFFFh specifies or indicates that the entire track is being described.

More than one logical block may be described by this address descriptor.

Table 138 defines the order of the fields used for sorting bytes from index format address descriptors if the command using the address descriptors specifies sorting.

Table 138 – Sorting order for bytes from index format address descriptors

Bit:	(MSB) 63	...	40	39	...	32	31	...	(LSB) 0
	CYLINDER NUMBER field			HEAD NUMBER field			BYTES FROM INDEX field		

6.2.7 Physical sector format address descriptor

A format type of 101b specifies the physical sector format address descriptor (see table 139). This address descriptor contains the location of a track or a sector (see 4.3.2).

Table 139 – Physical sector format address descriptor (101b)

Byte	Bit	7	6	5	4	3	2	1	0
0	(MSB)								
...		CYLINDER NUMBER							
2									(LSB)
3		HEAD NUMBER							
4	(MSB)								
...		SECTOR NUMBER							
7									(LSB)

The CYLINDER NUMBER field contains the cylinder number (see 4.3.2).

The HEAD NUMBER field contains the head number (see 4.3.2).

The SECTOR NUMBER field contains the sector number (see 4.3.2). A SECTOR NUMBER field set to FFFF_FFFFh specifies or indicates that the entire track is being described.

More than one logical block may be described by this address descriptor.

Table 140 defines the order of the fields used for sorting physical sector format address descriptors if the command using the address descriptors specifies sorting.

Table 140 – Sorting order for physical sector format address descriptors

Bit:	(MSB) 63	...	40	39	...	32	31	...	(LSB) 0
	CYLINDER NUMBER field			HEAD NUMBER field			SECTOR NUMBER field		

6.3 Diagnostic parameters

6.3.1 Diagnostic parameters overview

See table 141 for references to the pages and descriptors for diagnostic parameters used by direct access block devices.

The diagnostic pages and their corresponding page codes for direct access block devices are defined in table 141.

Table 141 – Diagnostic page codes for direct access block devices

Diagnostic page name	Page code	Reference
Diagnostic pages assigned by SPC-4	30h to 3Fh	SPC-4
Rebuild Assist Input diagnostic page	42h	6.3.2
Rebuild Assist Output diagnostic page		6.3.3
SCSI enclosure services diagnostic pages	01h to 2Fh	SES-2
Supported Diagnostic Page diagnostic page	00h	SPC-4
Translate Address Input diagnostic page	40h	6.3.4
Translate Address Output diagnostic page		6.3.5
Obsolete	41h	
Vendor Specific diagnostic pages	80h to FFh	
Reserved for this standard	43h to 7Fh	

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC 14776-323:2017

6.3.2 Rebuild Assist Input diagnostic page

An application client sends a RECEIVE DIAGNOSTIC RESULTS command to retrieve a Rebuild Assist Input diagnostic page (see table 142), which provides information about whether the rebuild assist mode (see 4.20) is enabled or not and a device server's rebuild assist mode capabilities.

Table 142 – Rebuild Assist Input diagnostic page

Byte	Bit	7	6	5	4	3	2	1	0
0		PAGE CODE (42h)							
1		Reserved							
2	(MSB)	PAGE LENGTH ($4 + (2 \times n)$)							
3									
4		Reserved							ENABLED
5		Reserved							
6		Reserved							
7		PHYSICAL ELEMENT LENGTH (n)							
8		DISABLED PHYSICAL ELEMENT MASK (if any)							
...									
7 + n									
8 + n		DISABLED PHYSICAL ELEMENT (if any)							
...									
7 + (2 × n)									

The PAGE CODE field and the PAGE LENGTH field are defined in SPC-4 and shall be set to the values defined in table 142.

An ENABLED bit set to one indicates that the rebuild assist mode is enabled. An ENABLED bit set to zero indicates that the rebuild assist mode is disabled.

The PHYSICAL ELEMENT LENGTH field indicates the length in bytes of the DISABLED PHYSICAL ELEMENT MASK field and the length in bytes of the DISABLED PHYSICAL ELEMENT field.

The bits in the DISABLED PHYSICAL ELEMENT MASK field indicate the bits in the DISABLED PHYSICAL ELEMENT field that are supported. Each bit set to one in the DISABLED PHYSICAL ELEMENT MASK field indicates that the corresponding bit in the DISABLED PHYSICAL ELEMENT field is supported and may be set to one in a Rebuild Assist Output diagnostic page sent with a SEND DIAGNOSTIC command.

The bits in the DISABLED PHYSICAL ELEMENT field indicate the physical elements that are disabled in this logical unit. Each bit set to one indicates that a physical element is disabled, and the device server shall report predicted read errors and predicted write errors for the associated group of LBAs.

6.3.3 Rebuild Assist Output diagnostic page

An application client sends a SEND DIAGNOSTIC command to send a Rebuild Assist Output diagnostic page (see table 143) that:

- a) enables or disables rebuild assist mode (see 4.20.2); and/or
- b) puts the logical unit in a simulated failure mode by disabling physical elements in conjunction with rebuild assist mode (see 4.20.5).

Table 143 – Rebuild Assist Output diagnostic page

Byte	Bit	7	6	5	4	3	2	1	0
0		PAGE CODE (42h)							
1		Reserved							
2	(MSB)	PAGE LENGTH (4 + (2 × n))							
3		(LSB)							
4		Reserved							
5		ENABLE							
6		Reserved							
7		PHYSICAL ELEMENT LENGTH (n)							
8		DISABLED PHYSICAL ELEMENT MASK (if any)							
...									
7 + n									
8 + n		DISABLE PHYSICAL ELEMENT (if any)							
...									
7 + (2 × n)									

The PAGE CODE field and the PAGE LENGTH field are defined in SPC-4 and shall be set to the values defined in table 143.

An ENABLE bit set to one specifies that, after all fields in this diagnostic page have been validated:

- a) a self-test of the physical elements in the logical unit may be performed; and
- b) rebuild assist mode is enabled.

An ENABLE bit set to zero specifies that:

- a) rebuild assist mode shall be disabled;
- b) the other fields in this page shall be ignored; and
- c) all physical elements shall be enabled.

The PHYSICAL ELEMENT LENGTH field shall be set to the same value that is returned in the PHYSICAL ELEMENT LENGTH field in the Rebuild Assist Input diagnostic page.

If the PHYSICAL ELEMENT LENGTH field is not set to the same value that is returned in the PHYSICAL ELEMENT LENGTH field in the Rebuild Assist Input diagnostic page, then the device server shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST.

The device server shall ignore the DISABLED PHYSICAL ELEMENT MASK field.

Each bit in the DISABLE PHYSICAL ELEMENT field specifies a physical element that shall be disabled. A bit set to one in the DISABLE PHYSICAL ELEMENT field specifies that the device server shall respond to read commands and write commands specifying LBAs associated with that physical element as if the associated LBAs have

predicted errors. A bit set to zero in the DISABLE PHYSICAL ELEMENT field specifies that the device server shall respond to read commands and write commands specifying LBAs associated with that physical element as if the associated LBAs do not have predicted errors. If the ENABLE bit is set to one, and the DISABLE PHYSICAL ELEMENT field specifies:

- any bits set to one that are not supported by the logical unit;
- all bits that are supported by the logical unit are set to one; or
- setting to zero any bits that are set to one,

then the device server shall terminate the command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST.

6.3.4 Translate Address Input diagnostic page

Table 144 defines the Translate Address Input diagnostic page sent by a device server in response to a RECEIVE DIAGNOSTIC RESULTS command after a Translate Address Output diagnostic page (see 6.3.4) has been sent by an application client with the SEND DIAGNOSTIC command. If a Translate Address Output diagnostic page has not yet been processed by the device server, the results of a RECEIVE DIAGNOSTIC RESULTS command requesting this diagnostic page are vendor specific.

Table 144 – Translate Address Input diagnostic page

Byte	Bit	7	6	5	4	3	2	1	0
0		PAGE CODE (40h)							
1		Reserved							
2	(MSB)	PAGE LENGTH (n - 3)							
3									
4		Reserved				SUPPLIED FORMAT			
5		RAREA	ALTSEC	ALTTRK	Reserved		TRANSLATED FORMAT		
Translated address list									
6	(MSB)	TRANSLATED ADDRESS 1 (if any)							
...									
13		(LSB)							
		⋮							
n - 7	(MSB)	TRANSLATED ADDRESS x (if any)							
...									
n		(LSB)							

The PAGE CODE field is defined in SPC-4 and shall be set to the value shown in table 144 for the Translate Address Input diagnostic page.

The PAGE LENGTH field is defined in SPC-4.

The SUPPLIED FORMAT field contains the value from the SUPPLIED FORMAT field in the previous Translate Address Output diagnostic page (see 6.3.5).

A reserved area (RAREA) bit set to zero indicates that no part of the translated address falls within a reserved area of the medium (e.g., speed tolerance gap, alternate sector, or vendor reserved area). A RAREA bit set to one indicates that all or part of the translated address falls within a reserved area of the medium. If the entire translated address falls within a reserved area, then the device server may not return a translated address.

An alternate sector (ALTSEC) bit set to zero indicates that no part of the translated address is located in an alternate sector of the medium or that the device server is unable to determine this information. An ALTSEC bit set to one indicates that the translated address is located in an alternate sector of the medium. If the device server is unable to determine if all or part of the translated address is located in an alternate sector, then the device server shall set this bit to zero.

An alternate track (ALTRK) bit set to zero indicates that no part of the translated address is located on an alternate track of the medium. An ALTRK bit set to one indicates that part or all of the translated address is located on an alternate track of the medium or the device server is unable to determine if all or part of the translated address is located on an alternate track.

The TRANSLATED FORMAT field contains the value from the TRANSLATE FORMAT field in the previous Translate Address Output diagnostic page (see 6.3.4).

The TRANSLATED ADDRESS field(s) contains the address descriptor (see 6.2) that the device server translated from the address descriptor supplied by the application client in the previous Translate Address Output diagnostic page (see 6.3.5). Each field shall be in the format (see 6.2) specified in the TRANSLATED FORMAT field. If the short block format address descriptor (see 6.2.2) is specified, then the first four bytes of the TRANSLATED ADDRESS field shall contain the short block format address descriptor and the last four bytes shall contain 0000_0000h.

If the returned data is in short block format (see 6.2.2), long block format (see 6.2.5), or physical sector format (see 6.2.7) and the ADDRESS TO TRANSLATE field in the previous Translate Address Output diagnostic page covers more than one address after it has been translated (e.g., because of multiple physical sectors within a single logical block or multiple logical blocks within a single physical sector), then the device server shall return all possible addresses that are contained in the area specified by the address to be translated. If the returned data is in bytes from index format (see 6.2.6), the device server shall return a pair of translated values for each of the possible addresses that are contained in the area specified by the ADDRESS TO TRANSLATE field in the previous Translate Address Output diagnostic page. Of the pair of translated values returned, the first indicates the starting location and the second the ending location of the area.

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC 14776-323:2017

6.3.5 Translate Address Output diagnostic page

The Translate Address diagnostic pages provides a method for an application client to have a device server translate an address descriptor (see 6.2) from one format to another. The address descriptor to be translated is sent to the device server in the Translate Address Output diagnostic page with a SEND DIAGNOSTIC command and the results are returned by the device server in the Translate Address Input diagnostic page sent in response to a RECEIVE DIAGNOSTIC RESULTS command.

Table 145 defines the format of the Translate Address Output diagnostic page sent with the SEND DIAGNOSTIC command. The translated address returned in the Translate Address Input diagnostic page is defined in 6.3.4.

Table 145 – Translate Address Output diagnostic page

Byte	Bit	7	6	5	4	3	2	1	0								
0		PAGE CODE (40h)															
1		Reserved															
2	(MSB)	PAGE LENGTH (000Ah)															
3										(LSB)							
4		Reserved				SUPPLIED FORMAT											
5		Reserved				TRANSLATE FORMAT											
6	(MSB)	ADDRESS TO TRANSLATE															
...																	
13										(LSB)							

The PAGE CODE field and PAGE LENGTH field are defined in SPC-4 and shall be set to the values shown in table 145 for the Translate Address Output diagnostic page.

The SUPPLIED FORMAT field specifies the format (see 6.2) of the ADDRESS TO TRANSLATE field. If the device server does not support the requested format, then the device server shall terminate the SEND DIAGNOSTIC command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST.

The TRANSLATE FORMAT field specifies the format (see 6.2) the device server shall use for the result of the address translation. If the device server does not support the specified format, then the device server shall terminate the SEND DIAGNOSTIC command with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN PARAMETER LIST.

The ADDRESS TO TRANSLATE field contains a single address descriptor that the application client is requesting the device server to translate. The format of this field depends on the value in the SUPPLIED FORMAT field. The formats are described in 6.2. If the short block format address descriptor is specified, then the first four bytes of the ADDRESS TO TRANSLATE field shall contain the short block format address descriptor and the last four bytes shall contain 0000_0000h.

6.4 Log parameters

6.4.1 Log parameters overview

6.4.1.1 Summary of log pages

See table 146 for references to the log pages and their corresponding page codes and subpage codes for direct access block devices. See SPC-4 for a detailed description of logging operations.

Table 146 – Log page codes and subpage codes for direct access block devices

Log page name	Page code ^a	Subpage code ^a	Reference
Application Client log page	0Fh	00h	SPC-4
ATA PASS-THROUGH Results	16h	00h	SAT-3
Background Scan Results log page	15h	00h	6.4.2
Buffer Over-Run/Under-Run log page	01h	00h	SPC-4
Format Status log page	08h	00h	6.4.3
Informational Exceptions log page	2Fh	00h	SPC-4
Last n Deferred Errors Or Asynchronous Events log page	0Bh	00h	SPC-4
Last n Error Events log page	07h	00h	SPC-4
Logical Block Provisioning log page	0Ch	00h	6.4.4
Non-Medium Error log page	06h	00h	SPC-4
Non-volatile Cache log page	17h	00h	6.4.5
Protocol-Specific Port log pages	18h	00h to FEh	SPC-4
Read Error Counters log page	03h	00h	SPC-4
Self-Test Results log page	10h	00h	SPC-4
Solid State Media log page	11h	00h	6.4.6
Start-Stop Cycle Counter log page	0Eh	00h	SPC-4
Supported Log Pages log page	00h	FFh	SPC-4
Supported Log Pages and Subpages log page	00h	00h	SPC-4
Supported Subpages	01h to 3Fh	FFh	SPC-4
Temperature log page	0Dh	00h	SPC-4
Verify Error Counters log page	05h	00h	SPC-4
Write Error Counters log page	02h	00h	SPC-4
Vendor specific	30h to 3Eh	00h to FEh	n/a

^a All page code and subpage code combinations not shown in this table are reserved.

6.4.1.2 Setting and resetting log parameters

In a LOG SELECT command (see SPC-4), an application client may specify that:

- a) all the parameters in a log page or pages are to be reset (i.e., the PCR bit set to one and the PARAMETER LIST LENGTH field is set to zero); or

- b) individual parameters in log page are to be changed to specified new values (i.e., the PCR bit is set to zero and the PARAMETER LIST LENGTH field is not set to zero).

The device server processing of LOG SELECT commands (see SPC-4) that request changes to individual log parameters or reset all log parameters depend on the log parameter that is being changed or reset, and is specified in the table that defines the log parameter using the keywords defined in table 147 (also see SPC-4).

Table 147 – Keywords for resetting or changing log parameters

Keyword	Device server processing when:	
	PCR bit is set to one ^a	PCR bit is set to zero ^b
Always	Reset the log parameter.	Change the log parameter.
Reset Only	Reset the log parameter.	If any changes are requested in the PARAMETER VALUE field of the log parameter, then:
Never	Do not reset the log parameter; see the LOG SELECT command in SPC-4 for description of possible error conditions.	a) terminate the command with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST, and the additional sense code set to INVALID FIELD IN PARAMETER LIST; and b) do not make any requested changes in any field in any log parameter in any log page
^a If the PCR bit is set to one, and the PARAMETER LIST LENGTH field is not set to zero, then the device server shall terminate the LOG SELECT command (see SPC-4). ^b If the PCR bit is set to zero, and the PARAMETER LIST LENGTH field is set to zero. then no log parameters are changed (see SPC-4).		

6.4.2 Background Scan log page

6.4.2.1 Background Scan log page overview

Using the format shown in table 149, the Background Scan log page reports information about:

- background pre-scan operations (see 4.24.2) and background medium scan operations (see 4.24.3); and
- any logical blocks where an error was detected during a background scan operation.

The parameter codes for the Background Scan log page are defined in table 148.

Table 148 – Background Scan log page parameter codes

Parameter code	Description	Resettable or Changeable ^a	Reference	Support required
0000h	Background Scan Status	Never	6.4.2.2	Mandatory
0001h to 0800h	Background Scan Results	Reset Only	6.4.2.3	Optional ^b
8000h to AFFFh	Vendor specific		n/a	Optional
All others	Reserved			
^a The keywords in this column – Always, Reset Only, and Never – are defined in 6.4.1.2. ^b If the Background Scan log page is supported, then at least one Background Scan Results log parameter shall be supported.				

The Background Scan log page has the format defined in table 149.

Table 149 – Background Scan log page

Byte	Bit	7	6	5	4	3	2	1	0
0		DS (1b)	SPF (0b)	PAGE CODE (15h)					
1		SUBPAGE CODE (00h)							
2	(MSB)	PAGE LENGTH (n - 3)							
3		(LSB)							
Background scan parameters									
4		Background scan parameter [first] (if any)							
...									
		⋮							
...		Background scan parameter [last] (if any)							
n									

The disable save (DS) bit, the subpage format (SPF) bit, the PAGE CODE field, the SUBPAGE CODE field, and the PAGE LENGTH field are described in SPC-4.

The DS bit, the SPF bit, the PAGE CODE field, and the SUBPAGE CODE field shall be set to the values shown in table 149 for the Background Scan log page.

If the device server processes a LOG SELECT command with the PCR bit set to one (see SPC-4), then the device server shall:

- a) not change the values in the Background Scan Status log parameter; and
- b) delete all Background Scan Results log parameters.

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC 14776-323:2017

6.4.2.2 Background Scan Status log parameter

The Background Scan Status log parameter for the Background Scan log page has the format defined in table 150.

Table 150 – Background Scan Status log parameter format

Byte	Bit	7	6	5	4	3	2	1	0
0	(MSB)	PARAMETER CODE (0000h)							
1		(LSB)							
2		Parameter control byte – binary format list log parameter (see SPC-4)							
		DU	Obsolete	TSD	ETC	TMC		FORMAT AND LINKING	
3		PARAMETER LENGTH (0Ch)							
4	(MSB)	ACCUMULATED POWER ON MINUTES							
...									
7									
8		Reserved							
9		BACKGROUND SCAN STATUS							
10	(MSB)	NUMBER OF BACKGROUND SCANS PERFORMED							
11									
12	(MSB)	BACKGROUND SCAN PROGRESS							
13									
14	(MSB)	NUMBER OF BACKGROUND MEDIUM SCANS PERFORMED							
15									

The PARAMETER CODE field is described in SPC-4 and shall be set to the value shown in table 150 for the Background Scan Status log parameter.

The DU bit, the TSD bit, the ETC bit, the TMC field, and the FORMAT AND LINKING field for the Background Scan Status log parameter shall be set for a binary format list log parameter as described in SPC-4.

The PARAMETER LENGTH field is described in SPC-4 and shall be set to the value shown in table 150 for the Background Scan Status log parameter.

The ACCUMULATED POWER ON MINUTES field indicates the number of minutes the device server has been powered on since manufacturing.

Table 151 defines the BACKGROUND SCAN STATUS field.

Table 151 – BACKGROUND SCAN STATUS field

Code	Description
00h	No background scan operation is active.
01h	A background medium scan operation is active.
02h	A background pre-scan operation is active.
03h	A background scan operation was halted due to a fatal error.
04h	A background scan operation was halted due to a vendor specific pattern of errors.
05h	A background scan operation was halted due to the medium being formatted without the PLIST.
06h	A background scan operation was halted due to a vendor specific cause.
07h	A background scan operation was halted due to the temperature being out of the allowed range.
08h	Background medium scan operations are enabled (i.e., the EN_BMS bit is set to one in the Background Control mode page (see 6.5.4)), and no background medium scan operation is active (i.e., the device server is waiting for Background Medium Scan Interval timer expiration before starting the next background medium scan operation).
09h	A background scan operation was halted due to the S_L_FULL bit being set to one in the Background Control mode page (see 6.5.4) and the background scan results list being full.
0Ah	A background pre-scan operation was halted due to the Background Pre-scan Time Limit timer expiring.
0Bh to FFh	Reserved

The NUMBER OF BACKGROUND SCANS PERFORMED field indicates the number of background scan operations (i.e., the total number of background pre-scan operations plus the number of background medium scan operations) that have been performed since the SCSI target device was shipped by the manufacturer.

The BACKGROUND SCAN PROGRESS field indicates the percent complete of a background scan operation in progress. The returned value is a numerator that has 65 536 (i.e., 1_0000h) as its denominator. If there is no background scan operation in progress (i.e., no background scan operation has been initiated since power on or the most recent background scan operation has completed), then the device server shall set the BACKGROUND SCAN PROGRESS field to 0000h.

The NUMBER OF BACKGROUND MEDIUM SCANS PERFORMED field indicates the number of background medium scan operations that have been performed since the SCSI target device was shipped by the manufacturer. If the NUMBER OF BACKGROUND MEDIUM SCANS PERFORMED field contains 0000h, then the number of background medium scan operations is unknown.

The total number of background pre-scan operations that have been performed is the value in the NUMBER OF BACKGROUND SCANS PERFORMED field minus the value in the NUMBER OF BACKGROUND MEDIUM SCANS PERFORMED field.

6.4.2.3 Background Scan Results log parameter

The Background Scan Results log parameter for the Background Scan log page has the format defined in table 152. If the Background Scan log page is reset, then all Background Scan Results log parameters are discarded. If no errors have occurred during a background scan or the Background Scan log page has been reset, then no Background Scan Results log parameters shall be present.

Table 152 – Background Scan Results log parameter format

Byte	Bit	7	6	5	4	3	2	1	0
0	(MSB)	PARAMETER CODE (0001h to 0800h)							
1		(LSB)							
2		Parameter control byte – binary format list log parameter (see SPC-4)							
		DU	Obsolete	TSD	ETC	TMC	FORMAT AND LINKING		
3		PARAMETER LENGTH (14h)							
4	(MSB)	ACCUMULATED POWER ON MINUTES							
...									
7									
8		REASSIGN STATUS				SENSE KEY			
9		ADDITIONAL SENSE CODE							
10		ADDITIONAL SENSE CODE QUALIFIER							
11		Vendor specific							
...									
15									
16	(MSB)	LOGICAL BLOCK ADDRESS							
...									
23									

The PARAMETER CODE field is described in SPC-4 and shall be set to a value from 0001h through 0800h in sequence as errors are discovered during a background scan operation. When all of the supported parameter code values have been used, and a new error is discovered during a background scan operation, the oldest Background Scan Results log parameter in the list (i.e., the Background Scan Results log parameter with the smallest value in the ACCUMULATED POWER ON MINUTES field) shall be discarded, and the PARAMETER CODE field in the Background Scan Results log parameter for the new defect shall be set to the parameter code value of the discarded Background Scan Results log parameter.

The DU bit, the TSD bit, the ETC bit, the TMC field, and the FORMAT AND LINKING field for a Background Scan Results log parameter shall be set for a binary format list log parameter as described in SPC-4.

The PARAMETER LENGTH field is defined in SPC-4 and shall be set to the value shown in table 152 for the Background Scan Results log parameter.

The ACCUMULATED POWER ON MINUTES field indicates the number of minutes that the device server has been powered on since manufacturing at the time the background scan error reported in the Background Scan Results log parameter occurred.

Table 153 defines the REASSIGN STATUS field.

Table 153 – REASSIGN STATUS field (part 1 of 2)

Code	Reason			
	LOWIR bit ^a		Original error ^b	Additional conditions
	0	1		
1h	Yes	Yes	Recovered or unrecovered	The LBA has not yet been reassigned. ^c
2h	Yes	No	Recovered	The device server performed automatic read reassignment for the LBA (i.e., performed a reassign operation for the LBA and a write operation with recovered logical block data). ^d
4h	Yes	Yes	Recovered	The device server's attempt to perform automatic read reassignment failed. The logical block may or may not now have an uncorrectable error. ^c
5h	Yes	No	Recovered	The error was corrected by the device server rewriting the logical block without performing a reassign operation.
6h	Yes	Yes	Recovered or unrecovered	Either: a) an application client caused automatic write reassignment for the LBA with a command performing a write operation; or b) the LBPRZ bit is set to one in the Logical Block Provisioning VPD page (see 6.6.4), and an application client caused an unmap operation for the LBA. ^c
7h	Yes	Yes	Recovered or unrecovered	Either: a) an application client caused a reassign operation for the LBA with a REASSIGN BLOCKS command; or b) the LBPRZ bit is set to zero in the Logical Block Provisioning VPD page (see 6.6.4), and an application client caused an unmap operation for the LBA. ^c
<p>Key:</p> <p>Yes = specifies that a Background Scan Results log parameter shall be generated for the error. No = specifies that a Background Scan Results log parameter shall not be generated for the error</p> <p>^a The LOWIR bit in the Background Control mode page (see 6.5.4). ^b Type of error detected while reading the logical block referenced by the LBA specified by the LOGICAL BLOCK ADDRESS field in the Background Scan Results log parameter (see 6.4.2.3) during a background scan operation. ^c The REASSIGN STATUS field in a given log parameter changes from 1h or 4h to 6h, 7h, or 8h when a reassign operation, write operation, or unmap operation on the LBA succeeds or when a reassign operation on the LBA fails. After the LBA is reassigned, any subsequent medium error occurring for the LBA is reported in a new log parameter with the same value in the LOGICAL BLOCK ADDRESS field as the value in the LOGICAL BLOCK ADDRESS field in the log parameter for the previous medium error for the LBA. ^d The ARRE bit in the Read-Write Error Recovery mode page (see 6.5.8) controls automatic read reassignment based on errors detected during all read medium operations, including those that are part of background scan operations.</p>				

Table 153 – REASSIGN STATUS field (part 2 of 2)

Code	Reason			
	LOWIR bit ^a		Original error ^b	Additional conditions
	0	1		
8h	Yes	Yes	Recovered or unrecovered	An application client's request for a reassign operation for the LBA with a REASSIGN BLOCKS command failed. The logical block referenced by the LBA may or may not still have an uncorrectable error.
All others	Reserved			
Key:				
Yes = specifies that a Background Scan Results log parameter shall be generated for the error.				
No = specifies that a Background Scan Results log parameter shall not be generated for the error				
^a The LOWIR bit in the Background Control mode page (see 6.5.4). ^b Type of error detected while reading the logical block referenced by the LBA specified by the LOGICAL BLOCK ADDRESS field in the Background Scan Results log parameter (see 6.4.2.3) during a background scan operation. ^c The REASSIGN STATUS field in a given log parameter changes from 1h or 4h to 6h, 7h, or 8h when a reassign operation, write operation, or unmap operation on the LBA succeeds or when a reassign operation on the LBA fails. After the LBA is reassigned, any subsequent medium error occurring for the LBA is reported in a new log parameter with the same value in the LOGICAL BLOCK ADDRESS field as the value in the LOGICAL BLOCK ADDRESS field in the log parameter for the previous medium error for the LBA. ^d The ARRE bit in the Read-Write Error Recovery mode page (see 6.5.8) controls automatic read reassignment based on errors detected during all read medium operations, including those that are part of background scan operations.				

If sense data is available, then the device server shall set the SENSE KEY field, the ADDITIONAL SENSE CODE field, and the ADDITIONAL SENSE CODE QUALIFIER field to a hierarchy of additional information relating to error conditions that occurred during the background scan operation. The content of these fields is represented in the same format used by the sense data (see SPC-4).

The LOGICAL BLOCK ADDRESS field indicates the LBA associated with the medium error.

6.4.3 Format Status log page

6.4.3.1 Format Status log page overview

Using the format shown table 155, the Format Status log page reports information about the most recent successful format operation and the state of the direct access block device since that operation was performed. The parameter codes for the Format Status log page are listed in table 154.

Table 154 – Format Status log page parameter codes

Parameter code	Description	Resettable or Changeable ^a	Reference	Support Required
0000h	Format Data Out	Never	6.4.3.2	Mandatory
0001h	Grown Defects During Certification	Never	6.4.3.3	Mandatory
0002h	Total Blocks Reassigned During Format	Never	6.4.3.4	Mandatory
0003h	Total New Blocks Reassigned	Never	6.4.3.5	Mandatory
0004h	Power On Minutes Since Format	Never	6.4.3.6	Mandatory
0005h to 7FFFh	Reserved			
8000h to FFFFh	Vendor specific			Optional
^a The keywords in this column – Always, Reset Only, and Never – are defined in 6.4.1.2.				

The Format Status log page has the format defined in table 155.

Table 155 – Format Status log page

Byte	Bit	7	6	5	4	3	2	1	0
0		DS (1b)	SPF (0b)	PAGE CODE (08h)					
1		SUBPAGE CODE (00h)							
2	(MSB)	PAGE LENGTH (n - 3)							
3		(LSB)							
Format Status log parameters									
4		Format Status log parameter [first]							
...		⋮							
...		Format Status log parameter [last]							
n									

The disable save (DS) bit, the subpage format (SPF) bit, the PAGE CODE field, the SUBPAGE CODE field, and the PAGE LENGTH field are described in SPC-4.

The DS bit, the SPF bit, the PAGE CODE field, and the SUBPAGE CODE field shall be set to the values shown in table 155 for the Format Status log page.

If a format operation has never been performed by the logical unit, then the log parameter for each Format Status log parameter listed in table 154 is not defined by this standard. If a device server begins a format operation, then the device server shall set each byte of the log parameter data (i.e., bytes four to n of the log parameter), if any, to FFh for each Format Status log parameter (e.g., if the PARAMETER LENGTH field is set to 02h, then the log parameter data is set to FFFFh).

If the most recent format operation failed or the information for a Format Status log parameter is not available, then the device server shall return FFh in each byte of the log parameter data (i.e., bytes four to n of the log parameter), if any, for the Format Status log parameter (e.g., if the PARAMETER LENGTH field is set to 04h, then the log parameter data shall be set to FFFF_FFFFh). The device server shall set each Format Status log parameter to be a multiple of four bytes.

6.4.3.2 Format Data Out log parameter

The Format Data Out log parameter of the Format Status log page has the format defined in table 156.

Table 156 – Format Data Out log parameter format

Byte	Bit	7	6	5	4	3	2	1	0
0	(MSB)	PARAMETER CODE (0000h)							
1		(LSB)							
2	Parameter control byte – binary format list log parameter (see SPC-4)								
	DU	Obsolete	TSD	ETC	TMC		FORMAT AND LINKING		
3		PARAMETER LENGTH (n - 3)							
4	(MSB)	FORMAT DATA OUT							
...									
n									

The PARAMETER CODE field is described in SPC-4 and shall be set to the value shown in table 156 for the Format Data Out log parameter.

The DU bit, the ETC bit, the TMC field, and the FORMAT AND LINKING field for the Format Data Out log parameter shall be set for a binary format list log parameter as described in SPC-4.

The target save disable (TSD) bit (see SPC-4) shall be set to zero for the Format Data Out log parameter, indicating that the logical unit saves the Format Data Out log parameter at vendor specific intervals without any request from an application client.

The PARAMETER LENGTH field is described in SPC-4.

After a successful format operation, the FORMAT DATA OUT field contains the FORMAT UNIT parameter list (see 5.3.2).

6.4.3.3 Grown Defects During Certification log parameter

The Grown Defects During Certification log parameter for the Format Status log page has the format defined in table 157.

Table 157 – Grown Defects During Certification log parameter format

Byte	Bit	7	6	5	4	3	2	1	0
0	(MSB)	PARAMETER CODE (0001h)							
1		(LSB)							
2	Parameter control byte – binary format list log parameter (see SPC-4)								
	DU	Obsolete	TSD	ETC	TMC		FORMAT AND LINKING		
3		PARAMETER LENGTH (08h)							
4	(MSB)	GROWN DEFECTS DURING CERTIFICATION							
...									
11									

The PARAMETER CODE field is described in SPC-4 and shall be set to the value shown in table 157 for the Grown Defects During Certification log parameter.

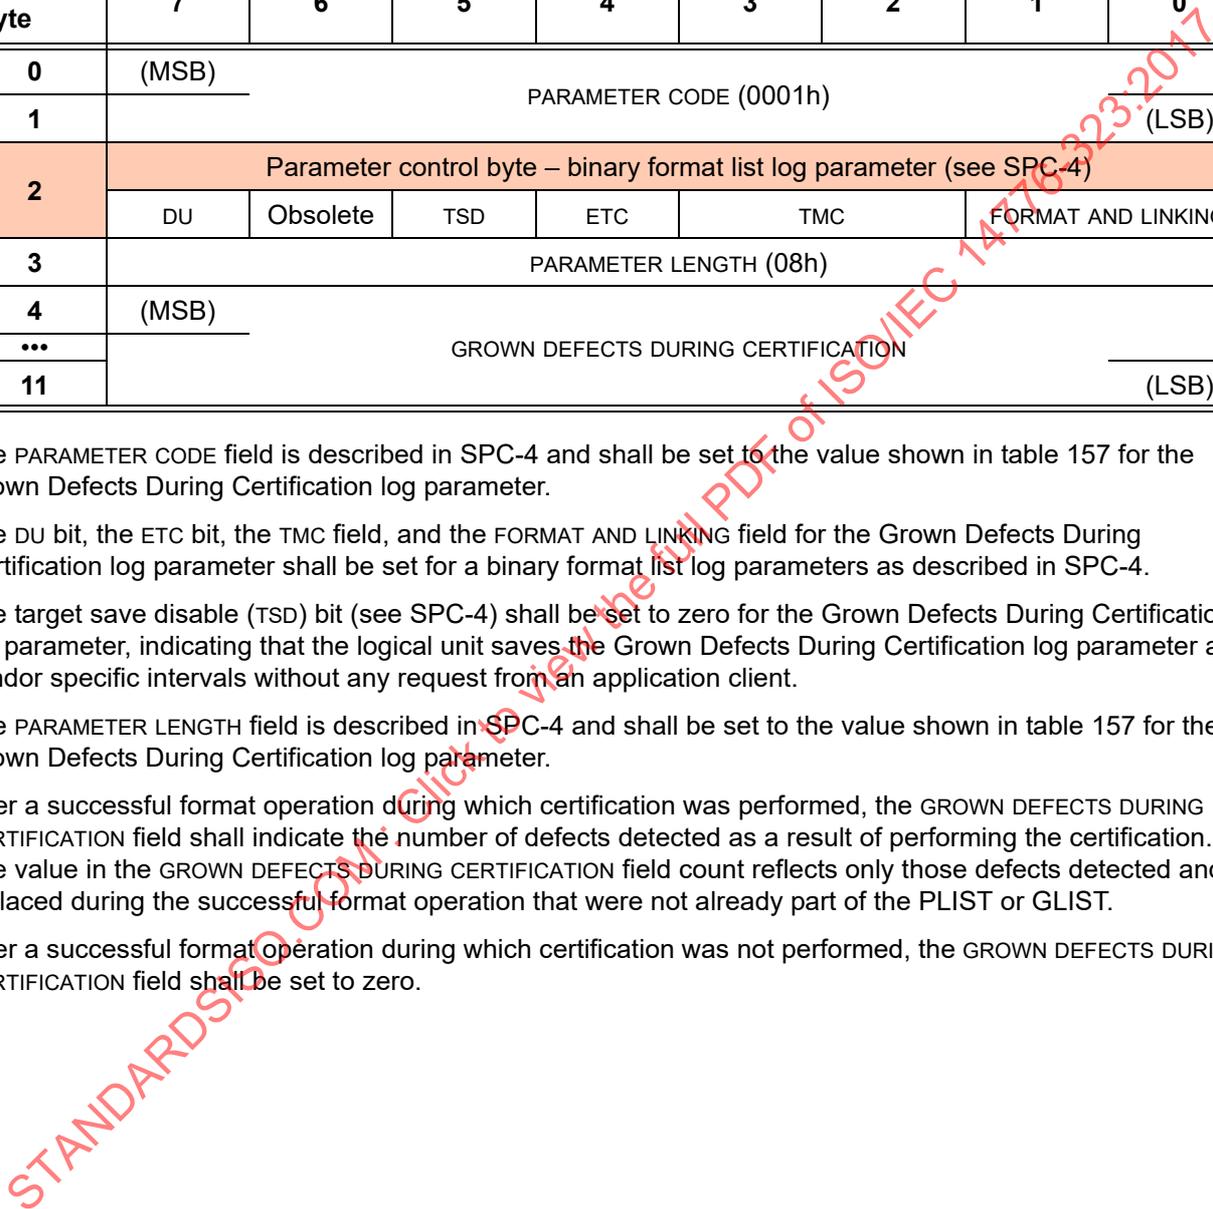
The DU bit, the ETC bit, the TMC field, and the FORMAT AND LINKING field for the Grown Defects During Certification log parameter shall be set for a binary format list log parameters as described in SPC-4.

The target save disable (TSD) bit (see SPC-4) shall be set to zero for the Grown Defects During Certification log parameter, indicating that the logical unit saves the Grown Defects During Certification log parameter at vendor specific intervals without any request from an application client.

The PARAMETER LENGTH field is described in SPC-4 and shall be set to the value shown in table 157 for the Grown Defects During Certification log parameter.

After a successful format operation during which certification was performed, the GROWN DEFECTS DURING CERTIFICATION field shall indicate the number of defects detected as a result of performing the certification. The value in the GROWN DEFECTS DURING CERTIFICATION field count reflects only those defects detected and replaced during the successful format operation that were not already part of the PLIST or GLIST.

After a successful format operation during which certification was not performed, the GROWN DEFECTS DURING CERTIFICATION field shall be set to zero.



6.4.3.4 Total Blocks Reassigned During Format log parameter

The Total Blocks Reassigned During Format log parameter for the Format Status log page has the format defined in table 158.

Table 158 – Total Blocks Reassigned During Format log parameter format

Byte	Bit	7	6	5	4	3	2	1	0
0	(MSB)	PARAMETER CODE (0002h)							
1		(LSB)							
2		Parameter control byte – binary format list log parameter (see SPC-4)							
		DU	Obsolete	TSD	ETC	TMC	FORMAT AND LINKING		
3		PARAMETER LENGTH (08h)							
4	(MSB)								
...		TOTAL BLOCKS REASSIGNED DURING FORMAT							
11		(LSB)							

The PARAMETER CODE field is described in SPC-4 and shall be set to the value shown in table 158 for the Total Blocks Reassigned During Format log parameter.

The DU bit, the ETC bit, the TMC field, and the FORMAT AND LINKING field for the Total Blocks Reassigned During Format log parameter shall be set for a binary format list log parameters described in SPC-4.

The target save disable (TSD) bit (see SPC-4) shall be set to zero for the Total Blocks Reassigned During Format log parameter, indicating that the logical unit saves the Total Blocks Reassigned During Format log parameter at vendor specific intervals without any request from an application client.

The PARAMETER LENGTH field is described in SPC-4 and shall be set to the value shown in table 158 for the Total Blocks Reassigned During Format log parameter.

The TOTAL BLOCKS REASSIGNED DURING FORMAT field contains the count of the total number of logical blocks that were reassigned during the most recent successful format operation.

6.4.3.5 Total New Blocks Reassigned log parameter

The Total New Blocks Reassigned log parameter for the Format Status log page has the format defined in table 159.

Table 159 – Total New Blocks Reassigned log parameter format

Byte	Bit	7	6	5	4	3	2	1	0
0	(MSB)	PARAMETER CODE (0003h)							
1		(LSB)							
2		Parameter control byte – binary format list log parameter (see SPC-4)							
		DU	Obsolete	TSD	ETC	TMC	FORMAT AND LINKING		
3		PARAMETER LENGTH (08h)							
4	(MSB)								
...		TOTAL NEW BLOCKS REASSIGNED							
11		(LSB)							

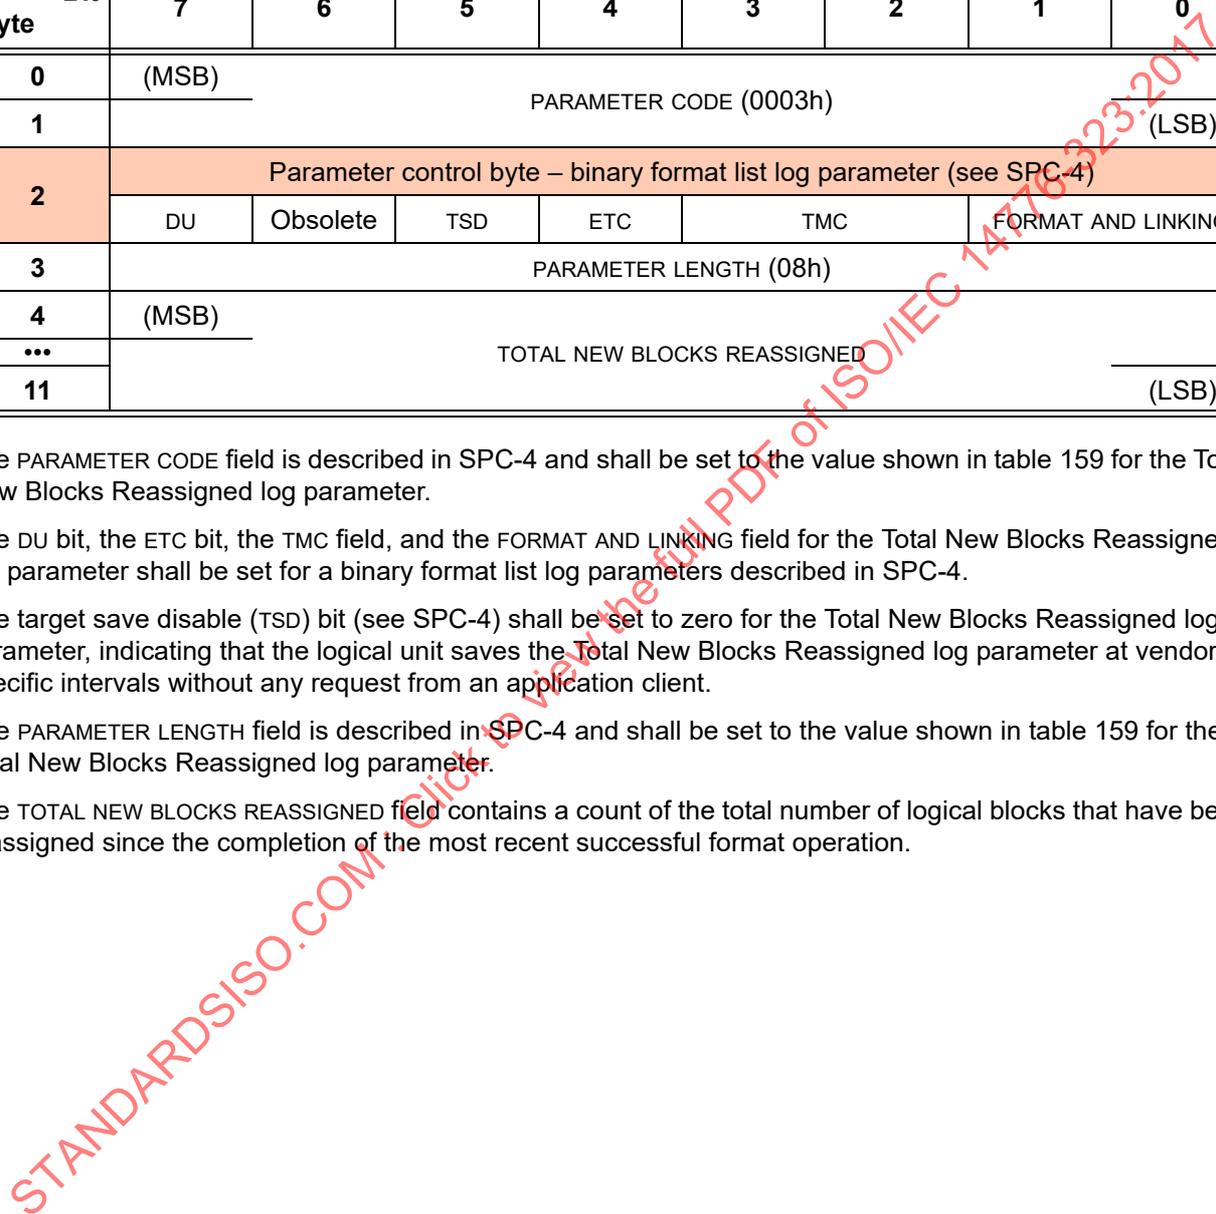
The PARAMETER CODE field is described in SPC-4 and shall be set to the value shown in table 159 for the Total New Blocks Reassigned log parameter.

The DU bit, the ETC bit, the TMC field, and the FORMAT AND LINKING field for the Total New Blocks Reassigned log parameter shall be set for a binary format list log parameters described in SPC-4.

The target save disable (TSD) bit (see SPC-4) shall be set to zero for the Total New Blocks Reassigned log parameter, indicating that the logical unit saves the Total New Blocks Reassigned log parameter at vendor specific intervals without any request from an application client.

The PARAMETER LENGTH field is described in SPC-4 and shall be set to the value shown in table 159 for the Total New Blocks Reassigned log parameter.

The TOTAL NEW BLOCKS REASSIGNED field contains a count of the total number of logical blocks that have been reassigned since the completion of the most recent successful format operation.



6.4.3.6 Power On Minutes Since Format log parameter

The Power On Minutes Since Format log parameter for the Format Status log page has the format defined in table 160.

Table 160 – Power On Minutes Since Format log parameter format

Byte	Bit	7	6	5	4	3	2	1	0
0	(MSB)	PARAMETER CODE (0004h)							
1		(LSB)							
2		Parameter control byte – binary format list log parameter (see SPC-4)							
		DU	Obsolete	TSD	ETC	TMC		FORMAT AND LINKING	
3		PARAMETER LENGTH (04h)							
4	(MSB)	POWER ON MINUTES SINCE FORMAT							
...									
7									

The PARAMETER CODE field is described in SPC-4 and shall be set to the value shown in table 160 for the Power On Minutes Since Format log parameter.

The DU bit, the ETC bit, the TMC field, and the FORMAT AND LINKING field for the Power On Minutes Since Format log parameter shall be set for a binary format list log parameter as described in SPC-4.

The target save disable (TSD) bit (see SPC-4) shall be set to zero for the Power On Minutes Since Format log parameter, indicating that the logical unit saves the Power On Minutes Since Format log parameter at vendor specific intervals without any request from an application client.

The PARAMETER LENGTH field is described in SPC-4 and shall be set to the value shown in table 160 for the Power On Minutes Since Format log parameter.

The POWER ON MINUTES SINCE FORMAT field contains the unsigned number of usage minutes (i.e., minutes with power applied regardless of power state) that have elapsed since the most recent successful format operation.

6.4.4 Logical Block Provisioning log page

6.4.4.1 Logical Block Provisioning log page overview

Using the format defined in table 162, the Logical Block Provisioning log page reports the logical block provisioning status of the logical unit. The parameter codes for the Logical Block Provisioning log page are listed in table 161.

Table 161 – Logical Block Provisioning log parameters

Parameter code ^a	Description	Resettable or Changeable ^b	Reference	Support Required
Resources that are associated with thresholds (0000h to 00FFh)				
0000h	Reserved			
0001h	Available LBA Mapping Resource Count	Never	6.4.4.2	Optional ^c
0002h	Used LBA Mapping Resource Count	Never	6.4.4.3	
0003h to 00FFh	Reserved			
Resources that are not associated with thresholds (0000h to 00FFh)				
0100h	De-duplicated LBA Resource Count	Never	6.4.4.4	Optional
0101h	Compressed LBA Resource Count	Never	6.4.4.5	
0102h	Total Efficiency LBA Resource Count	Never	6.4.4.6	
0103h to FFEFh	Reserved			
FFF0h to FFFFh	Vendor specific			
^a Parameter codes 0000h to 00FFh are coordinated with the THRESHOLD RESOURCE field in the threshold descriptor of the Logical Block Provisioning mode page (see 6.5.7). ^b The keywords in this column – Always, Reset Only, and Never – are defined in 6.4.1.2. ^c If this log page is supported, then at least one parameter shall be supported. A logical block provisioning log parameter in the range 0001h to 00FFh should be provided to report resource usage for each threshold resource for which a threshold descriptor in the Logical Block Provisioning mode page (see 6.5.7) is available.				

The Logical Block Provisioning log page has the format defined in table 162.

Table 162 – Logical Block Provisioning log page

Byte	Bit	7	6	5	4	3	2	1	0
0		DS (1b)	SPF (0b)	PAGE CODE (0Ch)					
1		SUBPAGE CODE (00h)							
2	(MSB)	PAGE LENGTH (n - 3)							
3		(LSB)							
Logical block provisioning parameter list									
4		Logical block provisioning log parameter [first]							
...									
		⋮							
...		Logical block provisioning log parameter [last]							
n									

The disable save (DS) bit, the subpage format (SPF) bit, the PAGE CODE field, the SUBPAGE CODE field, and the PAGE LENGTH field are described in SPC-4.

The DS bit, the SPF bit, the PAGE CODE field, and the SUBPAGE CODE field shall be set to the values shown in table 162 for the Logical Block Provisioning log page.

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC 14776-323:2017

6.4.4.2 Available LBA Mapping Resource Count log parameter

6.4.4.2.1 Available LBA Mapping Resource Count log parameter overview

The Available LBA Mapping Resource Count log parameter of the Logical Block Provisioning log page has the format defined in table 163.

Table 163 – Available LBA Mapping Resource Count log parameter format

Byte	Bit	7	6	5	4	3	2	1	0
0	(MSB)	PARAMETER CODE (0001h)							
1		(LSB)							
2		Parameter control byte – binary format list log parameter (see SPC-4)							
		DU	Obsolete	TSD	ETC	TMC	FORMAT AND LINKING		
3		PARAMETER LENGTH (08h)							
4	(MSB)	RESOURCE COUNT							
...									
7									
8		Reserved						SCOPE	
9		Reserved							
...									
11									

The PARAMETER CODE field is described in SPC-4 and shall be set to the value shown in table 163 for the Available LBA Mapping Resource Count log parameter.

The DU bit, the TSD bit, the ETC bit, the TMC field, and the FORMAT AND LINKING field for the Available LBA Mapping Resource Count shall be set for a binary format list log parameter as described in SPC-4.

The PARAMETER LENGTH field is described in SPC-4 and shall be set to the value shown in table 163 for the Available LBA Mapping Resource Count.

The RESOURCE COUNT field indicates an estimate of the number of available LBA mapping resources and is defined in 6.4.4.2.2.

The SCOPE field indicates the scope to which the RESOURCE COUNT field applies and is defined in table 164.

Table 164 – SCOPE field

Code	Description
00b	The scope of the resource count is not reported.
01b	The RESOURCE COUNT field indicates a resource that is dedicated to the logical unit. Usage of resources on other logical units does not impact the resource count.
10b	The resource count field indicates resources that may or may not be dedicated to any logical unit including the addressed logical unit. Usage of resources on other logical units may impact the resource count.
11b	Reserved

6.4.4.2.2 RESOURCE COUNT field

The RESOURCE COUNT field indicates an estimate of the number of LBA resources expressed as a number of threshold sets for the threshold resource indicated by the parameter code value. The nominal number of LBA resources is calculated as follows:

$$\text{LBA resources} = \text{resource count} \times \text{threshold set size}$$

where:

- resource count is the value in the RESOURCE COUNT field; and
- threshold set size is the number of LBAs in each threshold set (i.e., $2^{\text{(threshold exponent)}}$ LBAs, where the threshold exponent is indicated in the Logical Block Provisioning VPD page (see 6.6.4)).

6.4.4.3 Used LBA Mapping Resource Count log parameter

The Used LBA Mapping Resource Count log parameter of the Logical Block Provisioning log page has the format defined in table 165.

Table 165 – Used LBA Mapping Resource Count log parameter format

Byte	Bit	7	6	5	4	3	2	1	0
0	(MSB)	PARAMETER CODE (0002h)							
1		(LSB)							
2		Parameter control byte – binary format list log parameter (see SPC-4)							
		DU	Obsolete	TSD	ETC	TMC	FORMAT AND LINKING		
3		PARAMETER LENGTH (08h)							
4	(MSB)	RESOURCE COUNT							
...									
7									
8		Reserved					SCOPE		
9		Reserved							
...									
11									

The PARAMETER CODE field is described in SPC-4 and shall be set to the value shown in table 165 for the Used LBA Mapping Resource Count log parameter.

The DU bit, the TSD bit, the ETC bit, the TMC field, and the FORMAT AND LINKING field for the Used LBA Mapping Resource Count log parameter shall be set for a binary format list log parameter as described in SPC-4.

The PARAMETER LENGTH field is described in SPC-4 and shall be set to the value shown in table 165 for the Used LBA Mapping Resource Count log parameter.

The RESOURCE COUNT field indicates an estimate of the number of used LBA mapping resources and is defined in 6.4.4.2.2.

The SCOPE field indicates the scope to which the RESOURCE COUNT field applies and is defined in table 164.

6.4.4.4 De-duplicated LBA Resource Count log parameter

The De-duplicated LBA Resource Count log parameter of the Logical Block Provisioning log page (see table 166) contains information about de-duplicated LBA resources.

Table 166 – De-duplicated LBA Resource Count log parameter format

Byte	Bit	7	6	5	4	3	2	1	0
0	(MSB)	PARAMETER CODE (0100h)							
1		(LSB)							
2		Parameter control byte – binary format list log parameter (see SPC-4)							
		DU	Obsolete	TSD	ETC	TMC		FORMAT AND LINKING	
3		PARAMETER LENGTH (08h)							
4	(MSB)	RESOURCE COUNT							
...									
7									
8		Reserved						SCOPE	
9		Reserved							
...									
11									

The PARAMETER CODE field is described in SPC-4 and shall be set to the value shown in table 166 for the De-duplicated LBA Resource Count log parameter.

The DU bit, the TSD bit, the ETC bit, the TMC field, and the FORMAT AND LINKING field for the De-duplicated LBA Resource Count log parameter shall be set for a binary format list log parameter as described in SPC-4.

The PARAMETER LENGTH field is described in SPC-4 and shall be set to the value shown in table 166 for the De-duplicated LBA Resource Count log parameter.

The RESOURCE COUNT field indicates an estimate of the number of LBA resources made available as a result of de-duplication and is defined in 6.4.4.2.2.

The SCOPE field indicates the scope to which the RESOURCE COUNT field applies and is defined in table 164.

STANDARDSISO.COM: Click to view the full PDF of ISO/IEC 14776-323:2017

6.4.4.5 Compressed LBA Resource Count log parameter

The Compressed LBA Resource Count log parameter of the Logical Block Provisioning log page (see table 167) contains information about compressed LBA resources.

Table 167 – Compressed LBA Resource Count log parameter format

Byte	Bit	7	6	5	4	3	2	1	0
0	(MSB)	PARAMETER CODE (0101h)							
1		(LSB)							
2		Parameter control byte – binary format list log parameter (see SPC-4)							
		DU	Obsolete	TSD	ETC	TMC		FORMAT AND LINKING	
3		PARAMETER LENGTH (08h)							
4	(MSB)	RESOURCE COUNT							
...									
7									
8		Reserved						SCOPE	
9		Reserved							
...									
11									

The PARAMETER CODE field is described in SPC-4 and shall be set to the value shown in table 167 for the Compressed LBA Resource Count log parameter.

The DU bit, the TSD bit, the ETC bit, the TMC field, and the FORMAT AND LINKING field for the Compressed LBA Resource Count log parameter shall be set for a binary format list log parameter as described in SPC-4.

The PARAMETER LENGTH field is described in SPC-4 and shall be set to the value shown in table 167 for the Compressed LBA Resource Count log parameter.

The RESOURCE COUNT field indicates an estimate of the number of LBA resources made available as a result of compression and is defined in 6.4.4.2.2.

The SCOPE field indicates the scope to which the RESOURCE COUNT field applies and is defined in table 164.

6.4.4.6 Total Efficiency LBA Resource Count log parameter

The Total Efficiency LBA Resource Count log parameter of the Logical Block Provisioning log page (see table 168) contains information about the combined effects of all LBA resource efficiencies (e.g., the result of the combination of de-duplicated LBA resources and compressed LBA resources).

Table 168 – Total Efficiency LBA Resource Count log parameter format

Byte	Bit	7	6	5	4	3	2	1	0
0	(MSB)	PARAMETER CODE (0102h)							
1		(LSB)							
2		Parameter control byte – binary format list log parameter (see SPC-4)							
		DU	Obsolete	TSD	ETC	TMC	FORMAT AND LINKING		
3		PARAMETER LENGTH (08h)							
4	(MSB)	RESOURCE COUNT							
...									
7									
8		Reserved					SCOPE		
9		Reserved							
...									
11									

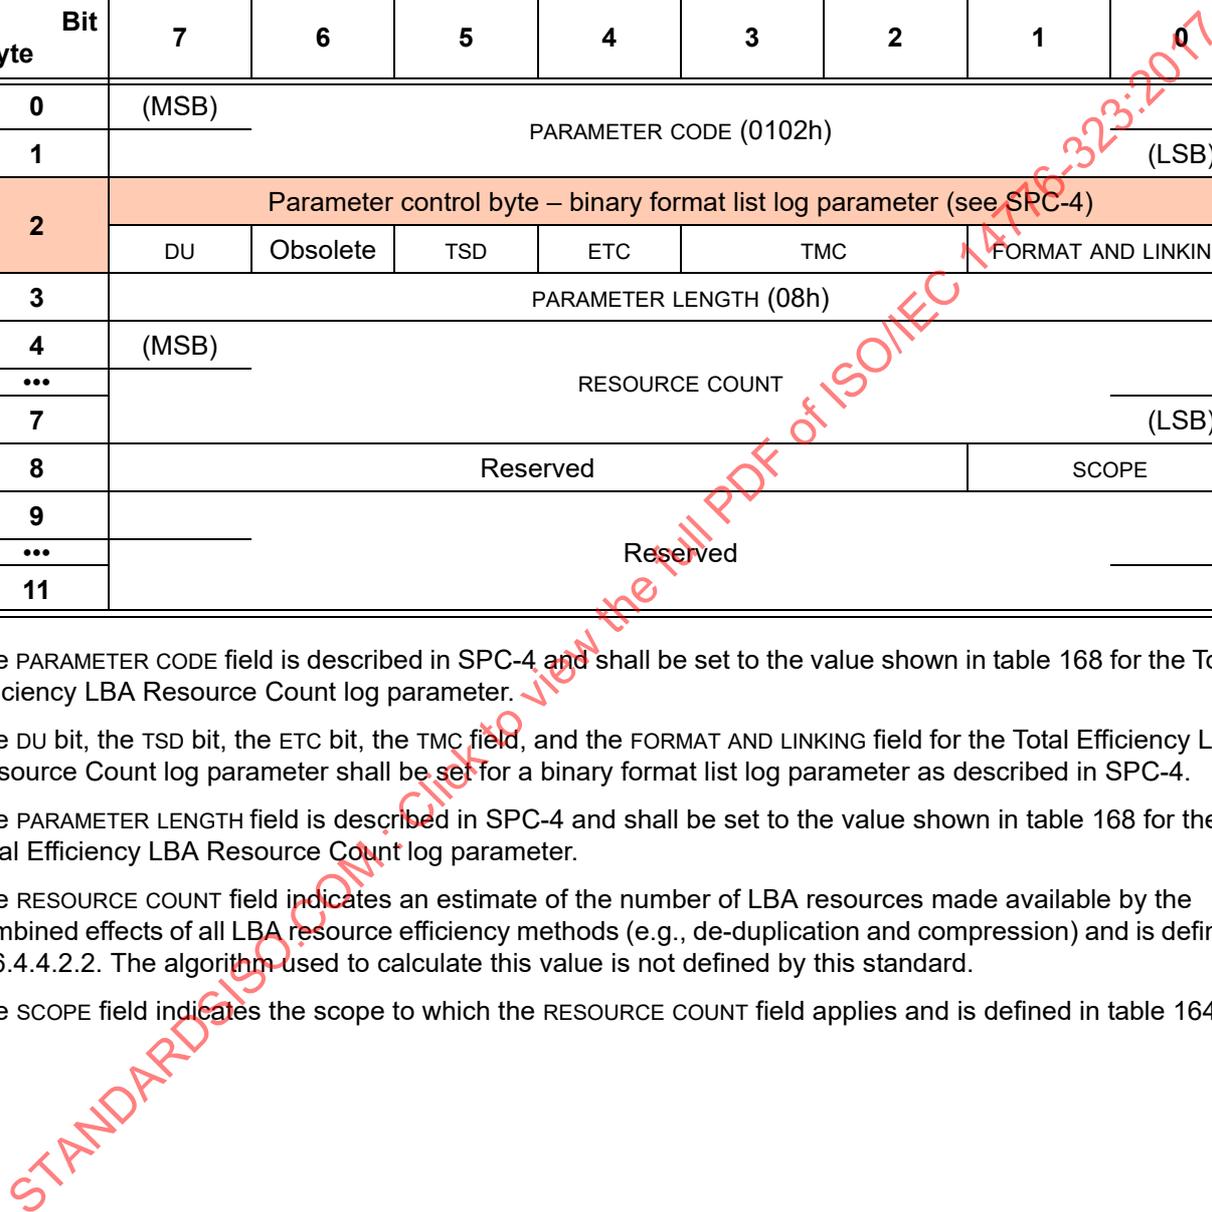
The PARAMETER CODE field is described in SPC-4 and shall be set to the value shown in table 168 for the Total Efficiency LBA Resource Count log parameter.

The DU bit, the TSD bit, the ETC bit, the TMC field, and the FORMAT AND LINKING field for the Total Efficiency LBA Resource Count log parameter shall be set for a binary format list log parameter as described in SPC-4.

The PARAMETER LENGTH field is described in SPC-4 and shall be set to the value shown in table 168 for the Total Efficiency LBA Resource Count log parameter.

The RESOURCE COUNT field indicates an estimate of the number of LBA resources made available by the combined effects of all LBA resource efficiency methods (e.g., de-duplication and compression) and is defined in 6.4.4.2.2. The algorithm used to calculate this value is not defined by this standard.

The SCOPE field indicates the scope to which the RESOURCE COUNT field applies and is defined in table 164.



6.4.5 Non-volatile Cache log page

6.4.5.1 Non-volatile Cache log page overview

Using the format shown in table 170, the Nonvolatile Cache log page reports the status of battery backup for a nonvolatile cache. The parameter codes for the Nonvolatile Cache log page are listed in table 169.

Table 169 – Nonvolatile Cache log parameters

Parameter code	Description	Resettable or Changeable ^a	Reference	Support Required
0000h	Remaining Nonvolatile Time	Never	6.4.5.2	Mandatory
0001h	Maximum Nonvolatile Time	Never	6.4.5.3	Mandatory
All others	Reserved			

^a The keywords in this column – Always, Reset Only, and Never – are defined in 6.4.1.2.

The Nonvolatile Cache log page has the format defined in table 170.

Table 170 – Nonvolatile Cache log page

Byte	Bit	7	6	5	4	3	2	1	0
0		DS	SPF (0b)	PAGE CODE (17h)					
1		SUBPAGE CODE (00h)							
2		(MSB)	PAGE LENGTH (n - 3)						
3									(LSB)
Nonvolatile cache log parameters									
4		Non-volatile cache log parameter [first] (see table 169)							
...									
		⋮							
...		Nonvolatile cache log parameter [last] (see table 169)							
n									

The disable save (DS) bit, the subpage format (SPF) bit, the PAGE CODE field, the SUBPAGE CODE field, and the PAGE LENGTH field are described in SPC-4.

The SPF bit, the PAGE CODE field, and the SUBPAGE CODE field shall be set to the values shown in table 170 for the Nonvolatile Cache log page.

6.4.5.2 Remaining Nonvolatile Time log parameter

The Remaining Nonvolatile Time log parameter of the Nonvolatile Cache log page has the format defined in table 171.

Table 171 – Remaining Nonvolatile Time parameter data

Byte	Bit	7	6	5	4	3	2	1	0	
0	(MSB)	PARAMETER CODE (0000h)							(LSB)	
1										
2		Parameter control byte – binary format list log parameter (see SPC-4)								
		DU	Obsolete	TSD	ETC	TMC		FORMAT AND LINKING		
3		PARAMETER LENGTH (04h)								
4		Obsolete								
5	(MSB)									
...		REMAINING NONVOLATILE TIME								
7									(LSB)	

The PARAMETER CODE field is described in SPC-4 and shall be set to the value shown in table 171 for the Remaining Nonvolatile Time log parameter.

The DU bit, the TSD bit, the ETC bit, the TMC field, and the FORMAT AND LINKING field for the Remaining Nonvolatile Time log parameter shall be set for a binary format list log parameter as described in SPC-4.

The PARAMETER LENGTH field is described in SPC-4 and shall be set to the value shown in table 171 for the Remaining Nonvolatile Time log parameter.

The REMAINING NONVOLATILE TIME field is defined in table 172.

Table 172 – REMAINING NONVOLATILE TIME field

Code	Description
00_0000h	Nonvolatile cache is volatile, either permanently or temporarily (e.g., if batteries require recharging).
00_0001h	Nonvolatile cache is expected to remain nonvolatile for an unknown amount of time (e.g., if battery status is unknown)
00_0002h to FF_FFFEh	Nonvolatile cache is expected to remain nonvolatile for the number of minutes indicated (e.g., for the life of the battery supplying power to random access memory).
FF_FFFFh	Nonvolatile cache is indefinitely nonvolatile.

6.4.5.3 Maximum Nonvolatile Time log parameter

The Maximum Nonvolatile Time log parameter of the Nonvolatile Cache log page has the format defined in table 173.

Table 173 – Maximum Nonvolatile Time parameter data

Byte	Bit	7	6	5	4	3	2	1	0	
0	(MSB)	PARAMETER CODE (0001h)							(LSB)	
1										
2		Parameter control byte – binary format list log parameter (see SPC-4)								
		DU	Obsolete	TSD	ETC	TMC		FORMAT AND LINKING		
3		PARAMETER LENGTH (04h)								
4		Obsolete								
5	(MSB)									
...		MAXIMUM NONVOLATILE TIME								
7									(LSB)	

The PARAMETER CODE field is described in SPC-4 and shall be set to the value shown in table 173 for the Maximum Nonvolatile Time log parameter.

The DU bit, the TSD bit, the ETC bit, the TMC field, and the FORMAT AND LINKING field for the Maximum Nonvolatile Time log parameter shall be set for a binary format list log parameter as described in SPC-4.

The PARAMETER LENGTH field is described in SPC-4 and shall be set to the value shown in table 173 for the Maximum Nonvolatile Time log parameter.

The MAXIMUM NONVOLATILE TIME field is defined in table 174.

Table 174 – MAXIMUM NONVOLATILE TIME field

Code	Description
00_0000h	Nonvolatile cache is volatile
00_0001h	Reserved
00_0002h to FF_FFFFh	Nonvolatile cache is capable of being nonvolatile for the estimated number of minutes indicated. If the time is based on batteries, then the time shall be based on the last full charge capacity rather than the design capacity of the batteries.
FF_FFFFh	Nonvolatile cache is indefinitely nonvolatile.

6.4.6 Solid State Media log page

6.4.6.1 Solid State Media log page overview

Using the format shown in table 176, the Solid State media log page reports parameters that are specific to SCSI target devices that contain solid state media. The parameter codes for the Solid State Media log page

are listed in table 175.

Table 175 – Solid State Media log parameters

Parameter code	Description	Resettable or Changeable ^a	Reference	Support Required
0001h	Percentage Used Endurance Indicator	Never	6.4.6.2	Mandatory
All others	Reserved			

^a The keywords in this column – Always, Reset Only, and Never – are defined in 6.4.1.2.

The Solid State Media log page has the format defined in table 176.

Table 176 – Solid State Media log page

Byte	Bit	7	6	5	4	3	2	1	0
0		DS	SPF (0b)	PAGE CODE (11h)					
1		SUBPAGE CODE (00h)							
2	(MSB)	PAGE LENGTH (n - 3)							
3		(LSB)							
Solid state media log parameters									
4		Solid state media parameter [first]							
...		⋮							
...		Solid state media parameter [last]							
n									

The disable save (DS) bit, the subpage format (SPF) bit, the PAGE CODE field, the SUBPAGE CODE field, and the PAGE LENGTH field are described in SPC-4.

The SPF bit, the PAGE CODE field, and the SUBPAGE CODE field shall be set to the values shown in table 176 for the Solid State Media log page.

6.4.6.2 Percentage Used Endurance Indicator log parameter

The Percentage Used Endurance Indicator log parameter of the Solid State Media log page has the format defined in table 177.

Table 177 – Percentage Used Endurance Indicator log parameter format

Byte	Bit	7	6	5	4	3	2	1	0
0	(MSB)	PARAMETER CODE (0001h)							
1		(LSB)							
2		Parameter control byte – binary format list log parameter (see SPC-4)							
		DU	Obsolete	TSD	ETC	TMC		FORMAT AND LINKING	
3		PARAMETER LENGTH (04h)							
4		Reserved							
...									
6									
7		PERCENTAGE USED ENDURANCE INDICATOR							

The PARAMETER CODE field is described in SPC-4 and shall be set to the value shown in table 177 for the Percentage Used Endurance Indicator log parameter.

The DU bit, the TSD bit, the ETC bit, the TMC field, and the FORMAT AND LINKING field for the Percentage Used Endurance Indicator log parameter shall be set for a binary format list log parameter as described in SPC-4.

The PARAMETER LENGTH field is described in SPC-4 and shall be set to the value shown in table 177 for the Percentage Used Endurance Indicator log parameter.

The PERCENTAGE USED ENDURANCE INDICATOR field indicates an estimate of the percentage of a SCSI target device that contains solid state media that has been used. The value in the field shall be set to zero at the time of manufacture. A value of 100 indicates that the estimated endurance of the SCSI target device that contain solid state media has been consumed, but may not indicate a SCSI target device that contain solid state media failure (e.g., minimum power-off data retention capability reached for SCSI target devices that contain solid state media using flash technology). The value is allowed to exceed 100. Values greater than 254 shall be reported as 255. The device server shall update the value at least once per power-on hour.