



Solid State Drive Specification

Ultrastar SSD400S

3.5" 4Gb FC-AL Solid State Drive

Models: HUSL4040ALF400
 HUSL4020ALF400
 HUSL4010ALF400

Version: 2.00

05 October 2012

Warning: Printed copies of this document are considered current only on the date of print. Replacement and disposal of down-level versions is the responsibility of the document holder.

1st Edition (Rev 2.00) (05 October 2012)

The following paragraph does not apply to the United Kingdom or any country where such provisions are inconsistent with local law: HGST, A WESTERN DIGITAL COMPANY, PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer or express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This publication could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. HGST may make improvements or changes in any products or programs described in this publication at any time.

It is possible that this publication may contain reference to, or information about, HGST products (machines and programs), programming, or services that are not announced in your country. Such references or information must not be construed to mean that HGST intends to announce such HGST products, programming, or services in your country.

Technical information about this product is available by contacting your local HGST representative or on the Internet at <http://www.hgst.com>

HGST may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents.

© Copyright HGST, a Western Digital company

Table of Contents

1.0 General	1
1.1 Introduction.....	1
1.2 Glossary	1
1.3 Caution.....	1
2.0 Outline of the Drive	3
3.0 Solid State Drive	5
3.1 Control Electronics	5
4.0 Drive Characteristics	7
4.1 Formatted Capacity.....	7
4.2 Data Sheet	7
4.3 Inquiry Information.....	8
4.3.1 Product ID.....	8
4.4 World Wide ID - Block Assignment	8
4.5 Performance characteristics	9
4.5.1 Drive ready time	9
4.5.2 SSD Command Overhead.....	9
4.5.3 SSD Response Time	9
4.5.4 Data transfer speed (Drive can sustain performance up to these values)	10
4.5.5 Random writes over the life of the drive	10
5.0 Data Integrity	11
5.1 Equipment Status	11
5.2 Error Recovery Procedure.....	11
6.0 Electrical Interface	13
6.1 FC-AL Connector	13
6.1.1 40 pin SCA-2 SFF-8045 FC-AL Connector Definition.....	13
6.1.2 Voltage and Ground Signals	14
6.1.3 Fault LED Out	14
6.1.4 Ready LED Out	14
6.1.5 Start Mated Controls.....	15
6.1.6 SEL_n and Enclosure Service Signals.....	15
6.1.7 -ENBL_BYN CH1, -ENBL_BYN_CH2.....	18
6.1.8 -Drive Present.....	18
6.1.9 Dev_Ctrl_Code_x	18
7.0 Environment	19
7.1 Temperature and humidity	19
7.2 Storage requirements	20
7.2.1 Packaging.....	20
7.2.2 Storage time	20
7.3 Corrosion test.....	20
7.4 Cooling requirements.....	21
8.0 DC Power Requirements	23
8.1 Power Supply Current, Average and Peak.....	23
8.2 Ripple Voltage	25
8.3 Power Consumption Efficiency Index	25

9.0 Reliability	27
9.1 Data Reliability	27
9.2 Failure prediction (S.M.A.R.T)	27
9.3 MTBF (Mean Time Between Failure): 2M hours.	27
9.4 Preventive Maintenance.....	27
9.5 Temperature Warning	27
10.0 Mechanical Specifications	29
10.1 Outline	29
10.2 Mechanical Dimensions.....	29
10.3 Mounting Positions and Tappings	31
10.4 Interface Connector.....	32
10.5 Drive Mounting.....	32
11.0 Acoustics, Vibration and Shock	33
11.1 Acoustics.....	33
11.2 Operating Vibration	33
11.2.1 Random Vibration.....	33
11.2.2 Swept Sine Vibration.....	33
11.3 Non-operating Vibrations	33
11.3.1 Random Vibration.....	33
11.3.2 Swept Sine Vibration.....	33
11.4 Operating shock	33
11.5 Non-operating shock.....	34
11.5.1 Half sinewave shock pulse.....	34
12.0 Identification	35
12.1 Labels.....	35
13.0 Electromagnetic Compatibility	37
13.1 Class A Regulatory Notices	38
14.0 Standards	41
14.1 UL and C-UL Standard Conformity	41
14.2 European Standards Compliance	41
14.3 German Safety Mark.....	41
14.4 Flammability.....	41
15.0 FC-AL attachment	43
15.1 Fundamentals	43
15.1.1 Node and Port names	44
15.1.2 NL_Port address	44
15.1.3 Primitive signals and sequences	45
15.1.4 Frames.....	46
15.1.5 Sequences.....	46
15.1.6 Exchanges	47
15.2 Basic Link Services	47
15.2.1 Abort sequence (ABTS).....	47
15.2.2 Basic accept (BA_ACC).....	48
15.2.3 Basic reject (BA_RJT).....	49
15.3 Extended Link Services	50
15.3.1 Link Service Accept (LS_ACC).....	51

15.3.2 Link Service Reject (LS_RJT).....	52
15.3.3 Port Login (PLOGI).....	53
15.3.4 Logout (LOGO).....	61
15.3.5 Fabric Login (FLOGI).....	62
15.3.6 Fabric Address Notification (FAN).....	66
15.3.7 Port Discovery (PDISC).....	67
15.3.8 Address Discovery (ADISC).....	67
15.3.9 Process Login (PRLI).....	68
15.3.10 Process Logout (PRLO).....	72
15.3.11 Read Link Error Status Block (RLS).....	74
15.3.12 Report Node Capabilities (RNC).....	75
15.3.13 Re-instate Recovery Qualifiers (RRQ).....	79
15.3.14 Third Party Process Logout (TPRLO).....	80
15.3.15 Request Node Identification Data (RNID).....	82
15.3.16 Report Port Speed Capabilities (RPSC).....	85
15.4 Common Fibre Channel Services.....	85
15.4.1 Register FC-4 Types (RFT_ID).....	88
15.5 FC-AL timers.....	90
15.5.1 Link Failure.....	91
15.6 Invalid frame delimiter.....	91
16.0 SCSI-FCP.....	93
16.1 Terminology.....	93
16.2 Information Units.....	94
16.2.1 FCP_CMND.....	95
16.2.2 FCP_XFER_RDY.....	97
16.2.3 FCP_DATA.....	97
16.2.4 FCP_RSP.....	97
16.3 Task Attributes.....	99
16.4 Task Management functions.....	100
16.4.1 Abort Task (Implemented as ABTS BLS).....	101
16.4.2 Abort Task Set.....	101
16.4.3 Terminate Task.....	102
16.4.4 Clear ACA.....	102
16.4.5 Target Reset.....	102
16.4.6 Clear Task Set.....	102
16.4.7 Reset LUN.....	102
16.5 Miscellaneous.....	102
16.5.1 Tags.....	102
16.5.2 Auto-Contingent Allegiance (ACA).....	102
16.5.3 Autosense.....	103
17.0 SCSI Command Set.....	105
17.1 SCSI Control Byte.....	107
17.2 Abbreviations.....	107
17.3 Byte ordering conventions.....	107
17.4 FORMAT UNIT (04).....	108
17.4.1 Parameter List Header.....	109

17.5 INQUIRY (12)	112
17.5.1 Inquiry Data	113
17.6 LOG SELECT (4C)	124
17.7 LOG SENSE (4D)	127
17.7.1 Log Page parameters	128
17.7.2 Log Sense Page 0	129
17.7.3 Log Sense Page 2	130
17.7.4 Log Sense Page 3	132
17.7.5 Log Sense Page 5	134
17.7.6 Log Sense Page 6	136
17.7.7 Log Sense Page D	137
17.7.8 Log Sense Page E	138
17.7.9 Log Sense Page F	140
17.7.10 Log Sense Page 10	141
17.7.11 Solid State Media log page (Page 11)	144
17.7.12 Log Sense Page 15	145
17.7.13 Log Sense Page 1A	147
17.7.14 Log Sense Page 2F	148
17.7.15 Log Sense Page 30	149
17.7.16 Log Sense Page 37	150
17.8 MODE SELECT (15)	152
17.9 MODE SELECT (55)	153
17.10 MODE SENSE (1A)	154
17.10.1 Mode Parameter List	155
17.10.2 Mode Page 00 (Vendor Unique Parameters)	159
17.10.3 Mode Page 01 (Read/Write Error Recovery Parameters)	161
17.10.4 Mode Page 02 (Disconnect/Reconnect Parameters)	162
17.10.5 Mode Page 03 (Format Device Parameters)	163
17.10.6 Mode Page 04 (Rigid Disk Drive Geometry Parameters)	165
17.10.7 Mode Page 07 (Verify Error Recovery Parameters)	166
17.10.8 Mode Page 08 (Caching Parameters)	167
17.10.9 Mode Page 0A (Control Mode Page Parameters)	169
17.10.10 Mode Page 0C (Notch Parameters)	171
17.10.11 Mode Page 19h (Fibre Channel Port Control Page)	172
17.10.12 Mode Page 1A (Power Control)	174
17.10.13 Mode Page 1C (Informational Exceptions Control)	175
17.11 MODE SENSE (5A)	178
17.12 PERSISTENT RESERVE IN (5E)	179
17.12.1 Service Action	179
17.12.2 Parameter data for Read Keys	180
17.12.3 Parameter Data for Read Reservations	181
17.13 PERSISTENT RESERVE OUT (5F)	182
17.13.1 Service Action	183
17.13.2 Type	183
17.13.3 Parameter list	184
17.13.4 Summary	185

17.14 PRE-FETCH (34)	187
17.15 READ (6) - (08)	188
17.16 READ (10) - (28)	189
17.17 READ (12) - (A8)	191
17.18 READ (16) - (88)	192
17.19 READ (32) - (7F/09)	193
17.20 READ BUFFER (3C)	195
17.20.1 Combined Header And Data (Mode 00000b)	196
17.20.2 Read Data (Mode 00010b)	196
17.20.3 Descriptor (Mode 00011b)	197
17.20.4 Read Data from Echo Buffer (Mode 01010b)	198
17.20.5 Echo Buffer Descriptor (Mode 01011b)	198
17.20.6 Expander Communications and Echo Buffer (Mode 11010b)	198
17.21 READ CAPACITY (10) - (25)	199
17.22 READ CAPACITY (16) (9E/10)	201
17.22.1 Returned Data Format	201
17.23 READ DEFECT DATA (37)	203
17.23.1 Defect List Header	205
17.23.2 Defect List Descriptor	205
17.23.3 Physical Sector Format (101b)	205
17.24 READ DEFECT DATA (B7)	207
17.24.1 Defect List Header	207
17.24.2 Defect List Descriptor	208
17.24.3 Physical Sector Format (101b)	208
17.25 READ LONG (3E)	209
17.26 REASSIGN BLOCKS (07)	210
17.27 RECEIVE DIAGNOSTICS RESULTS (1C)	212
17.27.1 Receive Diagnostic Results Page 0	212
17.27.2 Enclosure Service Information (ESI) Page Format	213
17.27.3 Receive Diagnostic Results Page 40	214
17.27.4 Receive Diagnostic Page A0h	215
17.28 RELEASE (17)	216
17.29 RELEASE (57)	217
17.30 REPORT DEVICE IDENTIFIER (A3/05)	218
17.31 REPORT LUNS (A0)	220
17.32 REPORT SUPPORTED OPERATION CODES (A3/0C)	221
17.32.1 All_commands parameter data format	222
17.32.2 One_command parameter data format	223
17.32.3 Command timeouts descriptor format	225
17.33 REPORT SUPPORTED TASK MANAGEMENT FUNCTIONS (A3/0D)	226
17.34 REQUEST SENSE (03)	227
17.35 RESERVE (16)	228
17.36 RESERVE (56)	229
17.37 REZERO UNIT (01)	230
17.38 SEEK (6) - (0B)	231
17.39 SEEK (10) - (2B)	231

17.40 SEND DIAGNOSTIC (1D)	232
17.40.1 Send Diagnostic Page 0	234
17.40.2 Send Diagnostic Page 40	234
17.40.3 Send Diagnostic Page A0h	236
17.41 SET DEVICE IDENTIFIER (A4/06)	237
17.42 START STOP UNIT (1B)	238
17.43 SYNCHRONIZE CACHE (10) - (35)	239
17.44 SYNCHRONIZE CACHE (16) - (91)	240
17.45 TEST UNIT READY (00)	241
17.46 UNMAP (42h)	242
17.46.1 UNMAP parameter list	242
17.47 VERIFY (2F)	244
17.48 VERIFY (12) - (AF)	246
17.49 VERIFY (16) - (8F)	247
17.50 VERIFY (32) - (7F/0A)	248
17.51 WRITE (6) - (0A)	250
17.52 WRITE (10) - (2A)	251
17.53 WRITE (12) - (AA)	253
17.54 WRITE (16) - (8A)	254
17.55 WRITE (32) - (7F/0B)	255
17.56 WRITE AND VERIFY (10) - (2E)	257
17.57 WRITE AND VERIFY (12) - (AE)	258
17.58 WRITE AND VERIFY (16) - (8E)	259
17.59 WRITE AND VERIFY (32) - (7F/0C)	260
17.60 WRITE BUFFER (3B)	261
17.60.1 Combined Header And Data (Mode 00000b)	261
17.60.2 Write Data (Mode 00010b)	262
17.60.3 Download Microcode (Mode 00100b)	262
17.60.4 Download Microcode and Save (Mode 00101b) -Single Binary File	263
17.60.5 Download Microcode and Save (Mode 00111b) - Multiple Binary Files	263
17.60.6 Write Data to Echo Buffer (Mode 01010b)	263
17.60.7 Enable Expander Communications Protocol (Mode 11010b)	263
17.61 WRITE LONG (3F)	264
17.62 WRITE SAME (41)	266
17.63 WRITE SAME (16) - (93)	267
17.64 WRITE SAME (32) - (7F/0D)	268
18.0 SCSI Status Byte	271
19.0 Additional information	273
19.1 Obtaining an AL_PA	273
19.2 Loop Initialization Procedure	273
19.3 Flow Control	275
19.4 Login Requirements	275
19.5 Public Loop Operation	276
19.5.1 NL_Port Initialization	277
19.6 SCSI Protocol	279
19.6.1 Priority of SCSI Status Byte Reporting	279

19.6.2 Invalid LUN Processing.....	279
19.6.3 Command Processing During Execution of Active I/O Process	279
19.6.4 Unit Attention Condition	281
19.6.5 Command Processing During Startup and Format Operations.....	283
19.6.6 Internal Error Condition.....	283
19.6.7 Deferred Error Condition.....	283
19.6.8 Degraded Mode.....	284
19.6.9 Command Processing while Reserved.....	290
19.7 Priority Commands	290
19.8 Command Queuing	291
19.8.1 Queue Depth	291
19.8.2 Queue Full Status.....	291
19.8.3 Termination of I/O Processes	291
19.9 Command Reordering.....	291
19.10 Concurrent I/O Process	291
19.11 Write Cache	291
19.12 Automatic Rewrite/Reallocate	292
19.13 Segmented Caching	293
19.13.1 Overview.....	293
19.13.2 Read Ahead.....	293
19.14 Multiple Initiator Systems	293
19.14.1 Sense Data.....	293
19.14.2 Mode Pages.....	293
19.15 Enclosure Services.....	293
19.15.1 Enclosure Initiated ESI	294
19.16 Multiple Initiator Environment.....	294
19.16.1 Initiator Sense Data.....	294
19.16.2 Initiator Mode Select/Mode Sense Parameters	294
19.17 Reset.....	294
19.17.1 Reset Sources	295
19.17.2 Reset Actions	295
19.18 Diagnostics.....	295
19.18.1 Power on Diagnostics	295
19.18.2 Self-test via SEND DIAGNOSTIC Command.....	296
19.19 Idle Time Function.....	299
19.20 Command Time out Limits	299
19.20.1 Reassignment Time.....	299
19.20.2 Format Time	299
19.20.3 Start/Stop Unit Time	299
19.20.4 Medium Access Command Time	300
19.20.5 Time-out Limits for Other Commands	300
19.21 Recommended Initiator ERP	301
19.21.1 Drive Service Strategy	301
19.21.2 Recommendations for System Error Log	302
19.21.3 Data Recovery Procedure	302
19.21.4 Nondata Error Recovery Procedure.....	302

20.0 SCSI Sense Data	309
20.1 SCSI Sense Data Format.....	309
20.2 Sense Data Description.....	310
20.2.1 Valid (Bit 7 of byte 0).....	310
20.2.2 Error Code (Bit 6 - 0 of byte 0)	310
20.2.3 ILI: Incorrect Length Indicator (Bit 5 of byte 2)	310
20.2.4 Sense Key (Bit 3 - 0 of byte 2)	311
20.2.5 Information Bytes (Byte 3 through 6).....	312
20.2.6 Additional Sense Length (Byte 7)	312
20.2.7 Command Specific Information (Byte 8 through 11)	312
20.2.8 Additional Sense Code/Qualifier (Byte 12 and 13).....	313
20.2.9 RU: Field Replaceable Unit (Byte 14).....	325
20.2.10 Sense Key Specific (Byte 15 through 17).....	325
20.2.11 Reserved (Byte 18 through 19).....	329
20.2.12 Vendor unique error information (Byte 20 through 23)	329
20.2.13 Physical Error Record (Byte 24 thru 29)	329
20.2.14 Reserved (Byte 30 through 31).....	329
21.0 Appendix. UEC list	331

List of Tables

Table 1.Product ID table.....	1
Table 2.Formatted Capacity.....	7
Table 3.Data Sheet.....	7
Table 4.Product ID in Inquiry Command.....	8
Table 5.Block assignment of World Wide ID in INQUIRY Command.....	8
Table 6.Drive ready time.....	9
Table 7.SSD Command Overhead.....	9
Table 8.SSD Response time.....	9
Table 9.40-pin Connector Signal Definition.....	13
Table 10.Start/Mated Controls.....	15
Table 11.SEL_ID/AL_PA.....	16
Table 12.SEL_N/P_ESI_N Signal Definition.....	17
Table 13.DEV_CTRL_CODE_N Signal Definition.....	18
Table 14.Operating and non-operating conditions.....	19
Table 15.Maximum allowable surface temperatures.....	21
Table 16.Input Voltage and Capacitance.....	23
Table 17.Power Supply Generated Ripple at Drive Power Connector.....	25
Table 18.Power Consumption Efficiency Index.....	25
Table 19.Physical Dimensions.....	29
Table 20.IEEE Registered Name format.....	44
Table 21.NL_Port address.....	44
Table 22.General frame format.....	46
Table 23.Frame header.....	46
Table 24.Basic link service command codes.....	47
Table 25.BA_ACC Payload.....	48
Table 26. BA_RJT Payload.....	49
Table 27.BA_RJT Reason Codes.....	49
Table 28.BA_RJT Reason Code Explanations.....	49
Table 29.Extended Link Service replies.....	50
Table 30.Extended Link Service requests.....	50
Table 31.Extended Link Service request 11h qualifiers.....	51
Table 32.LS_RJT payload.....	52
Table 33.LS_RJT reason codes.....	52
Table 34.LS_RJT reason code explanations.....	53
Table 35.PLOGI_REQ/PLOGI_ACC payload.....	54
Table 36.Common Service Parameter applicability (part 1 of 2).....	55
Table 37.Common Service Parameter applicability (part 2 of 2).....	56
Table 38.FC-PH Version.....	56
Table 39.Class Service Parameters.....	58
Table 40.LOGO payload.....	61
Table 41.LOGO ACC payload.....	61
Table 42.FLOGI_REQ/FLOGI_ACC payload.....	62
Table 43.Common Service Parameters (FLOGI_REQ/FLOGI_ACC).....	63

Table 44.Class Service Parameters	65
Table 45.FAN payload.....	66
Table 46.ADISC payload.....	67
Table 47.ADISC ACC payload	68
Table 48.PRLI payload.....	68
Table 49.PRLI ACC payload.....	69
Table 50.Login Service Parameter page.....	69
Table 51.Login Response Service Parameter page.....	71
Table 52.PRLI/PRLO ACC response codes	72
Table 53.PRLO payload	72
Table 54.PRLO ACC payload	73
Table 55.Logout Service Parameter page	73
Table 56.Logout Response Service Parameter page.....	74
Table 57.RLS payload	75
Table 58.RLS ACC payload	75
Table 59.Link Error Status block.....	75
Table 60.RNC/ACC payload	76
Table 61.RNC Capability Entry.....	77
Table 62.Document Identifiers	78
Table 63.RRQ payload	79
Table 64.RRQ ACC payload	79
Table 65.TPRLO payload.....	80
Table 66.TPRLO ACC payload.....	80
Table 67.Logout Service Parameter page.....	80
Table 68.Logout Response Service Parameter page.....	81
Table 69.RNID payload.....	82
Table 70.RNID accept payload.....	82
Table 71.Common Node Identification Data.....	83
Table 72.Topology Discovery Specific Node Identification Data.....	83
Table 73.Topology Discovery Unit Type	84
Table 74.Topology Discovery Flags.....	85
Table 75.RPSC payload.....	85
Table 76.RPSC ACC payload.....	85
Table 77.Payload of a CT Header.....	86
Table 78.Command/Response Codes	86
Table 79.RFT_ID payload	88
Table 80.RFT_ID ACC/RJT payload	88
Table 81.FS_RJT Reason Codes	89
Table 82.FS_RJT Reason Explanations.....	89
Table 83.FCAL timer values	90
Table 84.Information Units (IUs).....	94
Table 85.FCP_CMND payload.....	95
Table 86.FCP_CNTL field	95
Table 87.Task Attribute values.....	95
Table 88.TMF flag values.....	96
Table 89.FCP_XFER_RDY payload.....	97

Table 90.FCP_RSP payload	97
Table 91.FCP_STATUS field.....	98
Table 92.FCP_RSP_INFO field	99
Table 93.RSP_CODE definitions	99
Table 94.Task Management function RSP_CODE definitions	100
Table 95.SCSI Commands Supported	105
Table 96.SCSI Control Byte	107
Table 97.FORMAT UNIT (04)	108
Table 98.Format of the Parameter List Header.....	109
Table 99.Initialization Pattern Descriptor.....	110
Table 100.INQUIRY (12).....	112
Table 101.Page Code descriptions.....	112
Table 102.Inquiry Data- EVPD = 0	113
Table 103.Inquiry Data - EVPD = 1 (Page Code = 00h).....	115
Table 104.Inquiry Data - EVPD = 1 (Page Code = 03h).....	116
Table 105.Inquiry Data - EVPD = 1 (Page Code = 80h).....	118
Table 106.Inquiry Data - EVPD = 1 (Page Code = 83h).....	119
Table 107.Inquiry Data Format - EVPD = 1, (Page Code - 86h)	120
Table 108.Inquiry Data - EVPD = 1 (Page Code = B1h)	122
Table 109.Inquiry Data - EVPD = 1 (Page Code = D2h).....	123
Table 110.Log Select (4C).....	124
Table 111.Log Sense (4D)	127
Table 112.Log Sense Page 0.....	129
Table 113.Log Sense Page 2 (part 1 of 2).....	130
Table 114.Log Sense Page 2 (part 2 of 2).....	130
Table 115.Log Sense Page 3 (part 1 of 2).....	132
Table 116.Log Sense Page 3 (part 2 of 2).....	132
Table 117.Log Sense Page 5 (part 1 of 2).....	134
Table 118.Log Sense Page 5 (part 2 of 2).....	135
Table 119.Log Sense Page 6.....	136
Table 120.Log Sense Page D	137
Table 121.Log Sense Page E	138
Table 122.Log Sense Page F.....	140
Table 123.Log Sense Page F, Application Client Log	140
Table 124.Log Sense Page 10.....	141
Table 125.Log Sense Page 10, self-test results.....	141
Table 126.Log Sense Page 10, self-test results.....	142
Table 127.Log Sense Page 10, Extended Segment Number.....	143
Table 128.Solid State Media log page	144
Table 129.Solid State Media log page parameter codes	144
Table 130.Percentage Used Endurance Indicator parameter format	144
Table 131.Log Sense Page 15.....	145
Table 132.Log Sense Page 1A	147
Table 133. Log Sense Page 2F.....	148
Table 134.Vendor Unique parameter Code = 0000	148
Table 135.Log Sense Page 30.....	149

Table 136.Log Sense Page 37.....	150
Table 137.Mode Select (15).....	152
Table 138.Mode Select (55).....	153
Table 139.Mode Sense (1A).....	154
Table 140.Page Code Usage.....	155
Table 141.Mode parameter header (6).....	155
Table 142.Mode parameter header (10).....	156
Table 143.Mode Parameter Block Descriptor.....	157
Table 144.Mode Parameter Page Format.....	158
Table 145.Mode Parameter Page Format.....	158
Table 146.Vendor Unique Parameters - Page 00.....	159
Table 147.Mode Page 01 (Vendor Unique Parameters).....	161
Table 148.Mode Page 02 (Disconnect/Reconnect Parameters).....	162
Table 149.Mode Page 03 (Format Device Parameters).....	163
Table 150.Mode Page 04 (Rigid Disk Drive Geometry Parameters).....	165
Table 151.Mode Page 07 (Verify Error Recovery Parameters).....	166
Table 152.Page 08 (Caching Parameters).....	167
Table 153.Page 0A (Control Mode Page Parameters).....	169
Table 154.Control Extension Subpage.....	170
Table 155.Page 0C (Notch Parameters).....	171
Table 156.Mode Page 19h.....	172
Table 157.Values for RR_TOV Units.....	173
Table 158.Page 1A (Power Control).....	174
Table 159.Page 1C (Informational Exceptions Control).....	175
Table 160.Background Control (Subpage 01h).....	177
Table 161.Mode Sense (5A).....	178
Table 162.Persistent Reserve In (5E).....	179
Table 163.PERSISTENT RESERVE IN, Service Action Codes.....	179
Table 164.PERSISTENT RESERVE IN, parameter data for Read Keys.....	180
Table 165.PERSISTENT RESERVE IN, parameter data for Read Reservations.....	181
Table 166.PERSISTENT RESERVE IN, Read Reservation Descriptor.....	181
Table 167.PERSISTENT RESERVE OUT (5F).....	182
Table 168.PERSISTENT RESERVE OUT, Service Action Code.....	183
Table 169.PERSISTENT RESERVE OUT, Type Code.....	183
Table 170.Parameter List.....	184
Table 171.PERSISTENT RESERVE OUT, Service Action, Parameters.....	185
Table 172.APTPL and information held by a drive.....	186
Table 173.PRE-FETCH (34).....	187
Table 174.READ (6) - (08).....	188
Table 175.READ (10) - (28).....	189
Table 176.Read (12) - (A8).....	191
Table 177.READ (16) - (88).....	192
Table 178.READ (32) - (7F/09).....	193
Table 179.READ BUFFER (3C).....	195
Table 180.Read Buffer Header.....	196
Table 181.Read Buffer Description.....	197

Table 182.Echo Buffer Descriptor.....	198
Table 183.READ CAPACITY (10) - (25).....	199
Table 184.Format of READ CAPACITY command reply.....	200
Table 185.Read Capacity (16) (9E/10).....	201
Table 186.Returned Data Format.....	201
Table 187.P_TYPE field and PROT_EN bit.....	202
Table 188.LOGICAL BLOCKS PER PHYSICAL BLOCK EXPONENT field.....	202
Table 189.READ DEFECT DATA (37).....	203
Table 190.Defect List Format.....	204
Table 191.Defect List Header.....	205
Table 192.Defect List Descriptor.....	205
Table 193.Defect Descriptors of Physical Sector Format.....	205
Table 194.READ DEFECT DATA (B7).....	207
Table 195.Defect List Header.....	207
Table 196.Defect List Descriptor.....	208
Table 197.Defect Descriptors of Physical Sector Format.....	208
Table 198.READ LONG (3E).....	209
Table 199.REASSIGN BLOCKS (07).....	210
Table 200.Format of Reassign Blocks data.....	211
Table 201.RECEIVE DIAGNOSTIC RESULTS (1C).....	212
Table 202.Receive Diagnostic Results page 0.....	212
Table 203.Enclosure Page Support for Send and Receive Diagnostic Commands.....	213
Table 204.ESI Page Format.....	213
Table 205.Receive Diagnostic Results Page 40.....	214
Table 206.Translated address.....	214
Table 207.Device LED Control Page - Receive Diagnostic.....	215
Table 208.RELEASE (17).....	216
Table 209.RELEASE (57).....	217
Table 210.REPORT DEVICE IDENTIFIER (A3/05).....	218
Table 211.Report Device Identifier parameter list.....	219
Table 212.REPORT LUNS (A0).....	220
Table 213.LUN Reporting parameter list format.....	220
Table 214.REPORT SUPPORTED OPERATION CODES (A3/0C).....	221
Table 215.Reporting Options.....	221
Table 216.All_command parameter data format.....	222
Table 217.Command Descriptor format.....	222
Table 218. One_command parameter data format.....	223
Table 219.One_command parameter support field.....	224
Table 220.Command timeouts descriptor format.....	225
Table 221.Command timeouts descriptor Command Specific Field usage.....	225
Table 222.Report Supported Task Management Functions (A3/0D).....	226
Table 223.REQUEST SENSE (03).....	227
Table 224.RESERVE (16).....	228
Table 225.RESERVE (56).....	229
Table 226.REZERO UNIT (01).....	230
Table 227.SEEK (6) - (0B).....	231

Table 228. SEEK (10) - (2B).....	231
Table 229. SEND DIAGNOSTIC (1D).....	232
Table 230. SEND DIAGNOSTIC Function Code (1D).....	233
Table 231. Diagnostic Page 0.....	234
Table 232. Diagnostic Page 40.....	234
Table 233. Address to translate.....	235
Table 234. Device LED Control Page - Send Diagnostic.....	236
Table 235. SET DEVICE IDENTIFIER (A4/06).....	237
Table 236. SET DEVICE IDENTIFIER, Parameter List.....	237
Table 237. START STOP UNIT (1B).....	238
Table 238. SYNCHRONIZE CACHE (10) - (35).....	239
Table 239. Synchronize Cache (16) - (91).....	240
Table 240. TEST UNIT READY (00).....	241
Table 241. UNMAP (42h).....	242
Table 242. UNMAP parameter list.....	243
Table 243. UNMAP block descriptor.....	243
Table 244. VERIFY (2F).....	244
Table 245. Verify (12) - (AF).....	246
Table 246. Verify (16) - (8F).....	247
Table 247. Verify (32) - 7F/0A).....	248
Table 248. WRITE (6) - (0A).....	250
Table 249. WRITE (10) - (2A).....	251
Table 250. Write (12) - (AA).....	253
Table 251. Write (16) - (8A).....	254
Table 252. Write (32) - (7F/0B).....	255
Table 253. WRITE AND VERIFY (10) - (2E).....	257
Table 254. Write and Verify (12) - (AE).....	258
Table 255. Write and Verify (16) - (8E).....	259
Table 256. Write and Verify (32) - (7F/0C).....	260
Table 257. WRITE BUFFER (3B).....	261
Table 258. Write Buffer Header.....	262
Table 259. WRITE LONG (3F).....	264
Table 260. WRITE SAME (41).....	266
Table 261. Write Same (16) - (93).....	267
Table 262. ANCHOR bit, UNMAP Bit, and ANC_SUP bit.....	267
Table 263. Write Same (32) - (7F/0D).....	268
Table 264. SCSI Status Byte. Format of the SCSI STATUS byte.....	271
Table 265. Lixx Frame Payloads.....	274
Table 266. Response to Frames before PLOGI or PRLI.....	276
Table 267. Media Degraded Mode Becoming ready.....	285
Table 268. Media Degraded Mode - Context Load Failure.....	286
Table 269. Media Degraded Mode - Spindle Stopped by Unit Stop Command.....	287
Table 270. Self Configuration Failure Degraded Mode.....	288
Table 271. Format Command Failure Degraded Mode.....	289
Table 272. Short and Extended Self-Test Description.....	298
Table 273. Format of Sense Data.....	309

Table 274.Field Pointer Bytes	325
Table 275.Actual Retry Count	327
Table 276.Progress Indication	328
Table 277.Unit Error Codes.....	331

1.0 General

1.1 Introduction

This document describes the specifications of the following HGST 3.5 inch FCAL drives.

Table 1: Product ID table

Drive Name	Model Name	Type	Capacity (GB)	Interface
Ultrastar SSD400S-400	HUSSL4040ALF400	USSFCR400	400	3.5" FCAL
Ultrastar SSD400S-200	HUSSL4020ALF400	USSFCR200	200	3.5" FCAL
Ultrastar SSD400S-100	HUSSL4010ALF400	USSFCR100	100	3.5" FCAL

Note: The specifications in this document are subject to change without notice.

For technical and ordering information, please visit our website at <http://www.hgst.com>.

1.2 Glossary

Word	Meaning
BMS	Background Media Scan
Kb	Kilobit = 1000 bits
Mb	Megabit = 1,000,000 bits
GB	Gigabyte = 1,000,000,000 bytes
SSD	Solid State Drive
MB	Megabyte = 1,000,000 bytes
KB	Kilobyte = 1000 bytes
SFF	Small Form Factor
SMART	Self-Monitoring and Reporting Technology
FC-AL	Fibre Channel - Arbitrated Loop

1.3 Caution

This drive can be damaged by ESD (Electric Static Discharge). Any damage incurred to the drive after its removal from the shipping package and the ESD protective bag are the responsibility of the user.

2 Outline of the Drive

- Storage capacities of 400GB, 200 GB, and 100 GB
- 2 Gbps and 4 Gbps FCAL interface
- Supports dual-ported operations
- Supports full duplex operations
- Variable sector size (512B, 520B, and 528B) as production releases do not support 4096B
- Tagged Command Queuing support
- Automatic read/write data transfer
- Adaptive read ahead algorithm
- Write Cache via PLI protection
- XOR Function
- ECC On The Fly correction
- Automatic defect reallocation
- Self diagnostics at power on
- Use of SLC NAND Flash.
- SMART
- ANSI T10 Protection Information (End-to-End)

3.0 Solid State Drive

3.1 Control Electronics

The drive is electronically controlled by a microprocessor, logic modules, digital/analog modules and various drivers and receivers. The control electronics perform the following major functions:

- Monitors incoming power to insure safe writes
- Provides temporary back-up power in the event of a power loss
- Maintains data integrity through CRC, ECC and Power Loss Imminent detection

4.0 Drive Characteristics

4.1 Formatted Capacity

Table 2: Formatted Capacity

Description	HUSSL4040ALF400	HUSSL4020ALF400	HUSSL4010ALF400
Label capacity	400 GB	200 GB	100 GB
Total data bytes (512 bytes/sector)	400,088,457,216	200,049,647,616	100,030,242,816
Total logical data blocks	781,422,768 (2E9390B0h)	390,721,968 (1749F1B0h)	195,371,568 (BA52230h)

4.2 Data Sheet

Table 3: Data Sheet

Host Interface Transfer Rate	2 Gbps or 4.0 Gbps
Flash Media	34nm Single Level Cell NAND
SDRAM size	512 MB

4.3 Inquiry Information

4.3.1 Product ID

Product ID in Section 17.5.1.1, “Inquiry Data Format - EVPD = 0, Page Code = 0” on page 113, is as follows:

Table 4: Product ID in Inquiry Command

Product ID	Description
HUSSL4040ALF400	400 GB, FCAL
HUSSL4020ALF400	200 GB, FCAL
HUSSL4010ALF400	100 GB, FCAL

4.4 World Wide ID - Block Assignment

Block Assignment of World Wide ID is as follows:

Table 5: Block assignment of World Wide ID in INQUIRY Command

Manufacturing Site	Product	Block Assignment
China	HUSSL4040ALF400	001h ⁽¹⁾
	HUSSL4020ALF400	001h ⁽¹⁾
	HUSSL4010ALF400	001h ⁽¹⁾

Note (1) - Additional block assignment will be issued as needed based on actual production volume.

4.5 Performance characteristics

Drive performance is characterized by the following parameters:

- Command overhead
- Data transfer speed
- Buffering operation (read ahead/write cache)

Note: All the above parameters contribute to drive performance. There are other parameters that contribute to the performance of the actual system. This specification tries to define the bare drive characteristics, not system throughput, which depends on the system and the application.

4.5.1 Drive ready time

Table 6: Drive ready time

Model	
400/200/100 (GB)	<10 seconds to Data Ops, <1 second to Interface Response

4.5.2 SSD Command Overhead

Table 7: SSD Command Overhead

Model	
400/200/100 (GB)	30 usec

4.5.3 SSD Response Time

Table 8: SSD Response time

Model	Typical (Sec)	Max (Sec)
400/200/100 (GB)	100 Micro Second	20 Milli Second

4.5.4 Data transfer speed (Drive can sustain performance up to these values)

SSD400S	SLC FC
Random Read IOPS	
4KB, aligned, QD=1	7K
4KB, aligned, QD=4	22K
4KB, aligned, QD=32	40K
8KB, aligned, QD=1	6K
8KB, aligned, QD=4	18K
8KB, aligned, QD=32	23K
Random Write IOPS	
4KB, aligned, QD=1	8K
4KB, aligned, QD=4	16K
4KB, aligned, QD=32	20K
8KB, aligned, QD=1	6K
8KB, aligned, QD=4	12K
8KB, aligned, QD=32	12K
Random 70/30% Read/Write IOPS	
4KB, aligned, QD=1	8K
4KB, aligned, QD=4	19K
4KB, aligned, QD=32	24K
8KB, aligned, QD=1	6K
8KB, aligned, QD=4	14K
8KB, aligned, QD=32	16K
Sequential Read Data Transfer Rate (MB/s)	
64KB, aligned, QD=32	385
8KB, aligned, QD=32	16K

Notes:

Drive performance varies with model capacity and actual drive use

For this table, '1 MB / Sec' should be interpreted as **1024X1024** bytes per Second.

4.5.5 Random writes over the life of the drive

Model	100 GB	200GB	400GB
Worst Case (PB)	11	21.6	41.3

5.0 Data Integrity

The SSD employs a fail-safe write cache that insures customer data is committed to the media in the event of a power loss. This fail-safe write cache cannot be disabled. On a given write command, if all data has been received by the SSD and the RESPONSE frame has been sent to the Initiator with good status, it is guaranteed that all write data will be committed to the media in the event of a power loss. If the RESPONSE frame was not sent and some (or all) of the data was received by the SSD, some (or all) of the data may be committed to the media in the event of a power loss. In this case, it is guaranteed that all affected blocks will be readable without error, but not all blocks may return the newly written data.

5.1 Equipment Status

Equipment status is available to the host system any time the drive is not ready to READ or WRITE. This status normally exists at power-on time and will be maintained until the following conditions are satisfied:

- Self-check of drive is complete

Appropriate error status is made available to the host system if any of the following conditions occur after the drive has become ready:

- SMART error is detected

5.2 Error Recovery Procedure

Errors occurring with the drive are handled by the error recovery procedure.

Errors that are uncorrectable after application of the error recovery procedures are reported to the host system as non-recoverable errors.

6.0 Electrical Interface

6.1 FC-AL Connector

The drive conforms to SFF specifications 8045, 8067 and 8451. The SFF 8045/8067 documents defines the electrical specification and SFF-8451 defines the mechanical specification. The only difference between SFF-8045 and SFF-8067 is the Enclosure Service Interface. The drive will 'discover' the level of Enclosure Service Interface supported by the enclosure, and use the proper level for enclosure communication. This 'discovery' process is defined in the SFF-8067 specification.

The connector is expected to be used in an environment which uses a common connector structure for racking drives in a cabinet. The connector allows for plugging a drive directly into a backplane by providing the necessary electrical connection. Mechanical stability and device retention must be provided by a mechanism outside the drive.

The signals defined in this connector include:

- Dual ported Fibre Channel In/Out control signals
- +5V and +12V power
- ID select
- Motor start control
- LED drive signals
- PBC Interlock control
- Fibre Channel link rate selection/Hard Reset
- Power control

6.1.1 40 pin SCA-2 SFF-8045 FC-AL Connector Definition

The drive uses the connector compatible with the SFF-8045 connector specifications. The connector contact assignment is shown in following table.

Table 9: 40-pin Connector Signal Definition

Signal Name	Pin #	Pin #	Signal Name
-ENBL BYP CH1	1	21	12V Charge
12 Volts	2	22	12V Ground
12 Volts	3	23	12V Ground
12 Volts	4	24	+Port 1_In
-Parallel ESI	5*	25	-Port 1_In
-Drive Present	6	26	12V Ground
Ready LED Out	7	27	+Port 2_In
Power Control	8	28	-Port 2_In
Start_1/Mated	9	29	12V Ground
Start_2/Mated	10	30	+Port 1_Out
-ENBL BYP CH2	11	31	-Port 1_Out
Sel_6/-EFW	12*	32	5V Ground
Sel_5/-P_ESI_5	13*	33	+Port 2_Out

Sel_4/-P_ESI_4	14*	34	-Port 2_Out
Sel_3/-P_ESI_3	15*	35	5V Ground
Fault LED Out	16	36*	Sel_2/-P_ESI_2
Dev_Ctrl_Code_2	17	37*	Sel_1/-P_ESI_1
Dev_Ctrl_Code_1	18	38*	Sel_0/-P_ESI_0
5 Volts	19	39	Dev_Ctrl_Code_0
5 Volts	20	40	5V Charge

Note: The guide pins are connected to 5V ground.

* Definition changes for SFF-8067

6.1.2 Voltage and Ground Signals

The 12 V and 5 V contacts provide all of the voltages required by the drive. The two voltages share a common ground plane to which all of the ground contacts are connected.

6.1.3 Fault LED Out

As specified in the SFF-8045 specification, the Fault LED is driven under the following conditions:

- both enable bypass signals are asserted by the drive
- an internal failure has been detected by the drive
- the drive has been instructed by the host to turn on the LED

The drive provides an open-drain driver with up to 30mA of current sink capability to the drive fault LED. The cathode of the LED should be connected to this signal. The LED and the current-limiting resistor must be provided by the enclosure.

6.1.4 Ready LED Out

As specified in the SFF-8045 specification, the Ready LED Out signal has the following definition (per the Hot Plug Implementation):

- Drive not mated:
 - The signal is de-asserted (i.e. high). The LED is off.
- Drive mated, motor not spinning:
 - The signal is asserted (i.e. low) for a period long enough to be detected by the observer whenever a SCSI command is received. The LED is normally off.
- Drive mated, spinning up or down:
 - The signal is alternately asserted and de-asserted for a period of 1/2 second. The LED is flashing.
- Drive mated, motor spinning:
 - The signal is normally asserted continuously. The signal is de-asserted for a period long enough to be detected by an observer, whenever a SCSI command is received. The LED is usually on.

The drive provides an open-drain driver with up to 30mA of current sink capability to the Ready LED Out signal. The cathode of the LED should be connected to this signal. The LED and the current-limiting resistor must be provided by the enclosure.

6.1.5 Start Mated Controls

The Start_X Mated signals are TTL inputs to the drive and have 10K ohm pull-up resistors.

As per the SFF-8045 specification, the 'Start_x Mated' function is implemented as in the following table. Please refer to SFF-8045 for a more complete definition.

Table 10: Start/Mated Controls

Start_2 Mated	Start_1 Mated	Spin Function
Open	Open	Drive is not mated. No spin-up will occur.
Open	Gnd	The motor will spin up with the SCSI Unit Start Cmd
Gnd	Open	The motor will spin up after a delay of 6 times the modulo 8 of Sel_ID (in seconds)
Gnd	Gnd	The motor will spin up after drive initialization

6.1.6 SEL_n and Enclosure Service Signals

These signals have different definitions depending on the state of -Parallel ESI and the level of enclosure service supported by the backplane.

6.1.6.1 SEL_n and Enclosure Service Signals

-Parallel ESI is a TTL open-drain output from the drive. It is used to establish the definition of the SEL_n signals. When -Parallel ESI is de-asserted (high), the backplane shall preset SEL_ID information on these signals (within 1usec). When -Parallel ESI is asserted (low), the backplane (if supported) will present enclosure service information on these signals (within 1usec). The drive will then go through a 'discovery' phase to determine the level of enclosure services that the backplane supports (i.e. none, SFF-8045, or SFF-8067) and behave accordingly.

6.1.6.2 SEL_ID Function

The SEL_n (TTL compatible) inputs (defined when -Parallel ESI is de-asserted) provide a binary value of loop identifier to the drive. These 7 signals define 128 possible values and are directly translated into an 8 bit hard AL_PA via the table below. Only AL_PA's with neutral disparity are valid values. The drive will attempt to acquire this hard AL_PA for its own during the LIHA phase of the loop initialization process (LIP).

Table 11: SEL_ID/AL_PA

SEL_ID	AL_PA	SEL_ID	AL_PA	SEL_ID	AL_PA	SEL_ID	AL_PA
0	EF	20	B2	40	72	60	3A
1	E8	21	B1	41	71	61	39
2	E4	22	AE	42	6E	62	36
3	E2	23	AD	43	6D	63	35
4	E1	24	AC	44	6C	64	34
5	E0	25	AB	45	6B	65	33
6	DC	26	AA	46	6A	66	32
7	DA	27	A9	47	69	67	31
8	D9	28	A7	48	67	68	2E
9	D6	29	A6	49	66	69	2D
A	D5	2A	A5	4A	65	6A	2C
B	D4	2B	A3	4B	63	6B	2B
C	D3	2C	9F	4C	5C	6C	2A
D	D2	2D	9E	4D	5A	6D	29
E	D1	2E	9D	4E	59	6E	27
F	CE	2F	9B	4F	56	6F	26
10	CD	30	98	50	55	70	25
11	CC	31	97	51	54	71	23
12	CB	32	90	52	53	72	1F
13	CA	33	8F	53	52	73	1E
14	C9	34	88	54	51	74	1D
15	C7	35	84	55	4E	75	1B
16	C6	36	82	56	4D	76	18
17	C5	37	81	57	4C	77	17
18	C3	38	80	58	4B	78	10
19	BC	39	7C	59	4A	79	F
1A	BA	3A	7A	5A	49	7A	8
1B	B9	3B	79	5B	47	7B	4
1C	B6	3C	76	5C	46	7C	2
1D	B5	3D	75	5D	45	7E	1
1E	B4	3E	74	5E	43	7E	NA
1F	B3	3F	73	5F	3C	7F	-

Note: All values are in hex format. An AL_PA value of 0x00 is not valid for an NL_Port. A SEL_ID of 0x7F forces the drive to obtain a soft address.

6.1.6.3 P_ESI_n Function

The SEL_N (TTL compatible) I/O's (defined when -Parallel ESI is asserted) provide an interface between the enclosure and the drive. When signals are defined as outputs, they are configured as open-drain drivers with 4mA sink capability.

When the drive asserts -Parallel ESI, it goes through a 'discovery' process (see SFF-8067) to determine the level of enclosure service supported by the backplane. The outcome of the 'discovery' process is that the backplane:

1. Does not support enclosure services or
2. Supports SFF-8045 enclosure services or
3. Supports SFF-8067 enclosure services

The definition of the signals are summarized in the table below:

Table 12: SEL_N/P_ESI_N Signal Definition

-Parallel ESI	-Parallel ESI asserted (low)	
de-asserted (high)	SFF-8045*	SFF-8067*
Sel_0 (input)	P_ESI_0 (input)	Data(0) (bi-di**)
Sel_1 (input)	P_ESI_1 (input)	Data(1) (bi-di**)
Sel_2 (input)	P_ESI_2 (input)	Data(2) (bi-di**)
Sel_3 (input)	P_ESI_3 (input)	Data(3) (bi-di**)
Sel_4 (input)	P_ESI_4 (input)	-ENCL_ACK (output)
Sel_5 (input)	P_ESI_5 (input)	-DSK_RD (input)
Sel_6 (input)	-EFW (input)	-DSK_WR (input)

Note: * level of backplane support as established in 'discovery'
 ** the direction is determined by the state of DSK_RD and DSK_WR
 all signals are assumed valid 1 usec after -Parallel ESI changes state

6.1.6.4 SFF-8045 Enclosure Service Interface

The SFF-8045 Enclosure Service Interface defines 7 bits of enclosure status. This status is read by the drive, and presented to the initiator, upon receipt of the appropriate SCSI RECEIVE DIAGNOSTIC command. The definitions of the status bits are vendor specific. The drive does not try to interpret the status. The drive assumes -EFW is status and treats it as it does the other P_ESI_n signals.

6.1.6.5 SFF-8067 Enclosure Service Interface

The SFF-8067 Enclosure Service Interface defines a bidirectional communication path between the backplane and the drive. Read Communications are invoked by the drive upon receipt of the appropriate SCSI RECEIVE DIAGNOSTIC command. This has the system level effect of the host reading information from the enclosure. Write Communications are invoked by the drive upon receipt of the appropriate SCSI SEND DIAGNOSTIC command. This has the system level effect of the host writing information to the enclosure. The meaning of the information is vendor specific. The drive does not try to interpret the information. See the SFF-8067 Specification for a detailed description of the communication protocol.

6.1.7 –ENBL_BYP_CH1, -ENBL_BYP_CH2

These TTL outputs from the drive provide 4mA of sink capability. They are intended to control the state of a loop port bypass circuit on the backplane. The drive powers up with these signals turned off. It is assumed that the backplane will provide a 1Kohm pull-down resistor that will ensure the drive is bypassed on the loop when it is not present, or when it is powering up. After a successful power-up, the drive will attempt to enable itself on both loops (if allowed to do so via SCSI mode page 19h). These signals are also controllable by the host with the LPB and LPE fibre channel primitives.

Note: Contrary to industry standard, if the host attempts to bypass of either port1 or port2 with the LPB fibre channel primitive, the Ultrastar 15K600 generic configuration will assert the opposite ENBL-BYP signal. Please contact your HGST Customer Technical Support Representative, if any questions

6.1.8 –Drive Present

This signal is connected to the drive's ground plane. The backplane can use this signal to detect the presence of the drive.

6.1.9 Dev_Ctrl_Code_x

The signals DEV_CTRL_CODE_2 through DEV_CTRL_CODE_0 provide a binary code to the drive to control functions such as FC link rate, Power Failure Warning (PFW) and Hard Reset. The control function is either identified by a code or a sequence of codes on the DEV_CTRL_CODE signals. The table below defines the functions and assigned codes that use a decode of the value on the DEV_CTRL_CODE signals.

The Hard Reset function uses a sequence of values on the DEV_CTRL_CODE signals. The sequence is 5, 1, 3, 2, 3, 1, 5. A drive detecting a valid Hard Reset sequence shall perform the equivalent of a power-on-reset.

10 Kohm pull up resistors to 3.3 VOLTS are provided on the drive for both DEV_CTRL_CODE_2, DEV_CTRL_CODE_1 and DEV_CTRL_CODE_0 to be sure that each signal is maintained in its high state unless a low is provided from the backplane.

For more information refer to the SFF-8045 Specification.

Table 13: DEV_CTRL_CODE_N Signal Definition

Link Rate	DEV_CTRL_CODE_2	DEV_CTRL_CODE_1	DEV_CTRL_CODE_0
1.0625 GHz	1	1	1
2.1250 GHz	1	1	0
4.250 GHz	1	0	1
Reserved	1	0	0
Reserved	0	1	1
Reserved	0	1	0
Reserved	0	0	1
Power Failure Warning	0	0	0

7.0 Environment

7.1 Temperature and humidity

Table 14: <XREF>Operating and non-operating conditions

Operating conditions	
Ambient Temperature	0°C to 60°C
Relative humidity	5 to 90%, non-condensing
Maximum wet bulb temperature	29.4°C, non-condensing
Maximum surface temperature gradient	20 °C/hour
Altitude	-305 to 3,048 m
Shipping conditions	
Ambient Temperature	-55°C to 95°C
Relative humidity	5 to 95%, non-condensing
Maximum wet bulb temperature	35°C, non-condensing
Maximum surface temperature gradient	30°C/hour
Altitude	-305 to 12,192 m
Storage conditions	
Ambient Temperature	0°C to 60°C
Relative humidity	5 to 95%, non-condensing
Maximum wet bulb temperature	35°C, non-condensing
Altitude	-305 to 12,192 m

7.2 Storage requirements

7.2.1 Packaging

The drive or option kit is shipped in a sealed ESD bag by HGST.

7.2.2 Storage time

Cumulative storage time in the package must not exceed one year.

After the drive is unpackaged, it must not remain inoperative for longer than six months.

7.3 Corrosion test

The SSD shows no signs of corrosion inside or outside of the drive assembly and remains functional after being exposed to a temperature of 50°C and relative humidity of 90% for seven days.

7.4 Cooling requirements

Drive component temperatures must remain within the limits specified in the following table. Maximum component temperature ratings must not be exceeded under any operating condition. The drive may require forced air cooling to meet the specified, maximum operating temperatures.

Table 15: <XREF>Maximum allowable surface temperatures

Module name	Location	Maximum allowable surface temperature
SSD base	as noted in picture	70°C



Recommended Temperature Measurement Location

8.0 DC Power Requirements

The following voltage specification applies at the drive power connector. Connections to the drive should be made in a safety extra low voltage (SELV) circuit. There is no power on or power off sequencing requirement.

Adequate secondary over-current protection is the responsibility of the system.

Table 16: Input Voltage and Capacitance

Supply	Tolerance	Absolute Max Spike Voltage	Supply Rise Time	Capacitance
5 V	+/- 5%	5.5 V	0-200 ms	47uF
12 V	+/- 5%	15 V	0-400 ms	47 uF

8.1 Power Supply Current, Average and Peak

The following current and power requirements are typical when operating under the following conditions: Nominal 5 and 12V, Background Media Scan (BMS) disabled for Idle, Write Caching disabled and the drive reporting a temperature of 45C.

FCAL 100G	12V (A)	5V (A)	Power (W)	IOPS/MBPS
no power before test (Mean)	0.000	0.000	0.00	NA
no power before test (Max)	0.053	0.051		
Idle (Eean)	0.034	0.220	1.50	NA
Idle (max)	0.158	0.334		
standby (mean)	0.034	0.217	1.49	NA
standby (max)	0.152	0.339		
				NA
FCAL 100G	0.798	0.817		27542
FCAL 100G	0.060	0.367	2.56	
4k QD32 ranrw 70/30 (max)	0.340	0.859		
seqw (64K QD32) (mean)	0.147	0.415	3.84	335
seqw (64K QD32) (max)	0.459	0.922		
seqr (64K QD32) (mean)	0.066	0.390	2.74	388
seqr (64K QD32) (max)	0.199	0.714		

FCAL 200G	12V (A)	5V (A)	Power (W)	IOPS/MBPS
no power before test (Mean)	0.003	0.000	0.04	NA
no power before test (Max)	0.074	0.058		
Idle (mean)	0.030	0.227	1.50	NA
Idle (max)	0.163	0.339		
standby (mean)	0.031	0.224	1.49	NA
standby (max)	0.163	0.339		
				NA
start up (max)	0.788	0.880		24934
4k QD32 ranrw 70/30 (mean)	0.058	0.382	2.61	
4k QD32 ranrw 70/30 (max)	0.361	0.912		
seqw (64K QD32) (mean)	0.142	0.448	3.94	331
seqw (64K QD32) (max)	0.454	1.021		
seqr (64K QD32) (mean)	0.063	0.414	2.83	388
seqr (64K QD32) (max)	0.209	0.761		

FCAL 400G	12V (A)	5V (A)	Power (W)	IOPS/MBPS
no power before test (Mean)	0.005	0.000	0.06	NA
no power before test (Max)	0.069	0.058		
Idle (mean)	0.036	0.229	1.58	NA
Idle (max)	0.168	0.359		
standby (mean)	0.036	0.226	1.56	NA
standby (max)	0.168	0.323		
				NA
start up (max)	0.913	0.694		21621
4k QD32 ranrw 70/30 (mean)	0.076	0.395	2.89	
4k QD32 ranrw 70/30 (max)	0.402	1.120		
seqw (64K QD32) (mean)	0.168	0.488	4.45	338
seqw (64K QD32) (max)	0.433	1.005		
seqr (64K QD32) (mean)	0.083	0.437	3.17	388
seqr (64K QD32) (max)	0.225	0.766		

8.2 Ripple Voltage

Table 17: Power Supply Generated Ripple at Drive Power Connector

	Maximum (mV pp)	MHz
+5 V DC	250	0-10
+12 V DC	250	0-10

During drive start up, 12 volt ripple is generated by the drive (referred to as dynamic loading). If the power of several drives is daisy chained, the power supply ripple plus other drive dynamic loading must remain within the regulation tolerance of +5%. A common supply with separate power leads to each drive is a more desirable method of power distribution.

To prevent external electrical noise from interfering with the drive's performance, the drive must be held by four screws in a user system frame that has no electrical level difference at the four screw positions. The drive enclosure must not be used in the current return path of the drive power supply. The maximum common-mode noise current passing through the drive must not exceed 20 mA.

8.3 Power Consumption Efficiency Index

Table 18: Power Consumption Efficiency Index

Model	400 GB Model	200 GB Model	100 GB Model
Power Consumption Efficiency Index -Idle Mode (W/GB)	0.00395	0.0075	0.0015

9.0 Reliability

9.1 Data Reliability

- 22 bit correction per sector
- Offline full sector recovery through XOR for every N sectors (100G, N = 68: 200G, N = 136: 400G, N = 134)
- LBA seeded 32 bit CRC for ECC miscorrect detection
- Probability of uncorrectable data error is 1 in 1×10^{16} bits read

9.2 Failure prediction (S.M.A.R.T)

A recoverable equipment error is an error other than a seek/ID mismatch error or read error that is detected and corrected by the drive error recovery procedure. Examples are *Drive Not Ready* and internal drive errors.

SMART Monitoring Parameters are checked predict drive failure conditions before they occur. The primary parameters monitored for the SSD include:

1. Remaining Reserves: Ensures that the remaining spare erase blocks are at a sufficient level to guarantee proper operation of device.
2. Volatile Memory Backup: Self tests measure the capacitance of the power loss imminent circuitry to guarantee drive is able to commit data to media during unsafe power loss operations.
3. Wear Indicator: Endurance tracking mechanism based on maximum number of NAND erase operations performed on any band over the life of the device.

See “Log Sense Page 2F” on page 148 for tracking percentage of failure threshold for these parameters.

Non-recoverable equipment errors indicate a defective drive.

9.3 MTBF (Mean Time Between Failure): 2M hours.

This MTBF target is based on a sample population and is estimated by statistical measurements and acceleration algorithms under nominal operating conditions. MTBF ratings are not intended to predict an individual drive’s reliability. MTBF does not constitute a warranty.

9.4 Preventive Maintenance

None.

9.5 Temperature Warning

Temperature Warning is enabled by setting the EWASC (Enable Warning Additional Sense Code) bit to 1 and setting DEX-CPT (Disable Exception Control) bit to 0 in Mode Page 1C. For mode page settings, refer to Section “Mode Page 1C (Informational Exceptions Control)” on page 175. The warning is issued as sense data (Sense Key 01h, Code 0Bh, Qual 01h).

The drive temperature is reported in Log Sense page 2F. Refer to Section “Log Sense Page 2F” on page 148.

10.0 Mechanical Specifications



10.1 Outline




10.2 Mechanical Dimensions

Table 19: Physical Dimensions

Height [mm]	15.50 ± 0.25	
Width [mm]	101.60 ± 0.15	
Length (base) [mm]	146.80 ± 0.20	
Weight [grams - maximum]	400 GB Model	235 grams
	200 GB Model	226 grams
	100 GB Model	220 grams




10.3 Mounting Positions and Tappings



10.4 Interface Connector

The interface conforms to the specification SFF-8223, 2.5 Drive Form Factor with Serial Connector.



10.5 Drive Mounting

The drive will operate in all axes (6 directions). Performance and error rate will stay within specification limits if the drive is operated in the other orientations from which it was formatted.

The recommended mounting screw torque is 0.45 Nm (4.5 Kgf-cm). The recommended mounting screw depth is 2.5 mm maximum for bottom and 3.0 mm maximum for horizontal mounting.

Drive level vibration tests and shock tests are to be conducted with the drive mounted to a table using the bottom four screws.

11.0 Acoustics, Vibration and Shock

11.1 Acoustics

All SSD models have no acoustics, (0 bels).

11.2 Operating Vibration

11.2.1 Random Vibration

The drive is designed to operate without unrecoverable errors while being subjected to the vibration levels as defined below. The assessments are carried out during 30 minutes of random vibration using the power spectral density (PSD) levels as follows.

No Errors: 2.17 G RMS, 5-700 Hz, flat PSD profile for each of the three mutually perpendicular axes.

Note: The specified levels are measured at the mounting points.

11.2.2 Swept Sine Vibration

The drive will meet the criterion while operating in the respective conditions as described below.

No errors: 2.17 G RMS, 5-700 Hz.

11.3 Non-operating Vibrations

The drive will not sustain permanent damage or loss of recorded data after being subjected to the environments as described below.

11.3.1 Random Vibration

The test consists of a random vibration applied for each of the three mutually perpendicular axes. A time duration of ten minutes per axis.

3.13 G RMS, 5-800 Hz, flat PSD profile.

11.3.2 Swept Sine Vibration

The test consists of a swept sine vibration applied for each of the three mutually perpendicular axes.

3.13 G RMS, 10-800 Hz

11.4 Operating shock

The drive will meet the criterion while operating in the respective conditions as described below.

No data loss: 1000G, @0.5 ms duration, half sinewave shock pulse

500G, @ 2 ms duration, half sinewave shock pulse

The shock pulses of each level are applied to the drive, ten pulses for each direction and for all three mutually perpendicular axes. There must be a minimum of thirty seconds delay between shock pulses. The input level is applied to a base plate where the drive is attached using four mounting screws.

11.5 Non-operating shock

The drive will not sustain permanent damage or loss of recorded data after being subjected to the environments as described below.

11.5.1 Half sinewave shock pulse

100 G, 11 ms duration, half sinewave pulse

500 G, 2 ms duration, half sine wave pulse

1000 G, 0.5 ms duration, half sinewave pulse

The shocks are applied in each direction of the drive for the three mutually perpendicular axes, one axis at a time. The input level is applied to a base plate where the drive is attached using four mounting screws.

12.0 Identification

12.1 Labels

The following labels are affixed to every drive shipped from the drive manufacturing location in accordance with appropriate drive assembly drawing:

- A label containing HGST, a Western Digital Company logo, HGST part number and the statement “Made by HGST,” or HGST approved equivalent.
- A label containing drive model number, manufacturing date, formatted capacity, country of origin or HGST approved equivalent and UL, C-UL, TUV, CE, MIC, BSMI, CTICK, RoHS and Recycle logos.
- A bar code label containing the drive serial number.
- A user designed label, per agreement
- Interface definition mark, FCAL Model

The labels may be integrated with other labels.



Country of Origin: China

FCAL

17AUG2010



D33373



UL US
E240353 SG

HGS-HUSSL4040ALF40(A)
WARRANTY VOID IF ANY LABEL/
SCREW IS REMOVED OR BROKEN



MODEL: HUSSL4040ALF400	LES
FCAL	DC
RATED: 5V 0.5A 12V 0.2A	
CAPACITY: 400 GB	
P/N : 0B24940	F/W: XXXX



1MYYY123456789012



P/N: 0B24940



S/N: XQV02N0B

TYPE: USSFCR400

13.0 Electromagnetic Compatibility

The drive, when installed in a suitable enclosure and exercised with a random accessing routine at a maximum data rate will comply with the worldwide EMC requirements listed below.

The drive is designed for system integration and installation into a suitable enclosure for use. As such, the drive is supplied as a subassembly and is not subject to Subpart A of Part 15 of the FCC Rules and Regulations.

The design of the drive serves to minimize radiated emissions when installed in an enclosure that provides reasonable shielding. As such, the drive is capable of meeting FCC Class A limits. However, it is the users responsibility to assure that the drive meets the appropriate EMC requirements in their system. Shielded I/O cables may be required if the enclosure does not provide adequate shielding, with the shields grounded to the enclosure and to the host computer.

Radiated and Conducted EMI

CISPR22	Class A
AS/NZS CISPR22	Class A
CNS13438 (Taiwan)	Class A
EN55022 (EU)	Class A
FCC Title47 Part 15 (United States)	Class A
GB9254 (China)	Class A
ICES-003, Issue 4	Class A
VCCI (Japan)	Class A

ITE Immunity

EN55024

Power Line Harmonics

EN61000-3-2 (EU)

GB17625.1 (China)

Voltage Fluctuations and Flicker

EN61000-3-3 (EU)

GB17625.2 (China)

13.1 Class A Regulatory Notices

European Union

This product is in conformity with the protection requirements of EU Council Directive 2004/108/EC on the approximation of the laws of the Member States relating to electromagnetic compatibility. Conformity is based on compliance to the following harmonized standards:

- EN 55022: 2006 + A1:2007 (Class A)
- EN 55024: 1998 +A1:2001 +A2:2003
- EN 61000-3-2: 2006
- EN 61000-3-3:1995 + A1:2001 + A2:2005

This product is also in conformity with the protection requirements of EU Council Directive 2006/95/EC on the harmonization of the laws of the Member States relating to electrical equipment designed for use within certain voltage limits. Conformity is based on compliance to the following harmonized standards:

- EN 60950-1:2006 with Am. 11, Second Edition
- IEC 60950-1:2005, Second Edition
- UL 60950-1, Second Edition, 2007-03-27
- CSA C22.2 No. 60950-1-07, Second Edition, 2007-03

HGST cannot accept responsibility for any failure to satisfy the protection requirements resulting from a non-recommended modification of the product, including the fitting of non-HGST option cards.

This product has been tested and found to comply with the limits for Class A Information Technology Equipment according to European Standard EN 55022. The limits for Class A equipment were derived for commercial and industrial environments to provide reasonable protection against interference with licensed communication equipment.

Warning

This is a Class A product. In a domestic environment this product may cause radio interference in which case the user may be required to take adequate measures.

Canada

This Class A digital apparatus complies with Canadian ICES-003.

Cet appareil numérique de la classe A est conforme à la norme NMB-003 du Canada.

Germany

Deutschsprachiger EU Hinweis:

Hinweis für Geräte der Klasse A EU-Richtlinie zur Elektromagnetischen Verträglichkeit Dieses Produkt entspricht den Schutzanforderungen der EU-Richtlinie 89/336/EWG zur Angleichung der Rechtsvorschriften über die elektromagnetische Verträglichkeit in den EU-Mitgliedsstaaten und hält die Grenzwerte der EN 55022 Klasse A ein. Um dieses sicherzustellen, sind die Geräte wie in den Handbüchern beschrieben zu installieren und zu betreiben. Des Weiteren dürfen auch nur von der HGST empfohlene Kabel angeschlossen werden. HGST übernimmt keine Verantwortung für die Einhaltung der Schutzanforderungen, wenn das Produkt ohne Zustimmung der HGST verändert bzw. wenn Erweiterungskomponenten von Fremdherstellern ohne Empfehlung der HGST gesteckt/eingebaut werden. EN 55022 Klasse A Geräte müssen mit folgendem Warnhinweis versehen werden:

"Warnung: Dieses ist eine Einrichtung der Klasse A. Diese Einrichtung kann im Wohnbereich Funk-Störungen verursachen; in diesem Fall kann vom Betreiber verlangt werden, angemessene Maßnahmen zu ergreifen und dafür aufzukommen."

Deutschland: Einhaltung des Gesetzes über die elektromagnetische Verträglichkeit von Geräten Dieses Produkt entspricht dem "Gesetz über die elektromagnetische Verträglichkeit von Geräten (EMVG)". Dies ist die Umsetzung der EU-Richtlinie

89/336/EWG in der Bundesrepublik Deutschland.

Zulassungsbescheinigung laut dem Deutschen Gesetz über die elektromagnetische Verträglichkeit von Geräten (EMVG) vom 18. September 1998 (bzw. der EMC EG Richtlinie 89/336) für Geräte der Klasse A Dieses Gerät ist berechtigt, in Übereinstimmung mit dem Deutschen EMVG das EG-Konformitätszeichen - CE - zu führen. Verantwortlich für die Konformitätserklärung nach Paragraf 5 des EMVG ist die HGST, a Western Digital Company , 5600 Cottle road, San Jose, California 95193.

Informationen in Hinsicht EMVG Paragraf 4 Abs. (1) 4:

Das Gerät erfüllt die Schutzanforderungen nach EN 55024 und EN 55022 Klasse A.

Korea (MIC)

이 기기는 업무용으로 전자파적합등록을 한 기기이오니 판매자 또는 사용자는 이점을 주의하시기 바라며, 만약 잘못 판매 또는 구입하였을 때에는 가정용으로 교환하시기 바랍니다.

Taiwan (BSMI)

警告使用者：
這是甲類的資訊產品，在居住的環境中使用時，可能會造成射頻干擾，在這種情況下，使用者會被要求採取某些適當的對策。

新加坡商日立環球儲存科技股份有限公司台灣分公司
台北市敦化北路 167 號 5 樓 (宏國大樓)

14.0 Standards

The following shows the safety standards for different countries.

14.1 UL and C-UL Standard Conformity

The drive is certified under the following safety standards for use in Information Technology Equipment, including Electrical Business Equipment:

EN 60950-1:2006 with Am. 11, Second Edition, Europe

IEC 60950-1:2005, Second Edition, International

UL 60950-1, Second Edition, 2007-03-27, USA

CSA C22.2 No. 60950-1-07, Second Edition, 2007-03, Canada

The UL recognition, or the C-UL certification, is maintained for the duration of the product manufacturing life cycle. The UL and C-UL recognition marks appear on the drive label.

14.2 European Standards Compliance

This product is certified to the EN 60950-1:2006 with Am. 11, Second Edition safety standard for Europe.

14.3 German Safety Mark

The product is certified by TUV to meet EN 60950-1:2006 with Am. 11, Second Edition safety standard under the Bauart Mark.

14.4 Flammability

The printed wiring boards, and connectors used in this drive meet or exceed the UL minimum flammability classifications listed in the table below.

The flammability ratings are marked on the printed wiring boards and flex cables.

Component	Flammability Rating
Rigid Printed Wiring Board	Min. V-1
3.5" FCAL Connector	Min. V-2

15.0 FC-AL attachment

This section defines some basic terminology and describes the behavior of the drive when attached to a Fibre Channel Arbitrated Loop.

15.1 Fundamentals

This section introduces some of the terminology that is used in describing Fibre Channel and FC-AL.

Fibre Channel is logically a bi-directional serial data channel between two Nodes. Nodes are physically connected by a Link; the point of connection between the link and the node is called a Port.

Ports may be connected point-to-point by a single link or by a switching network (Fabric). If the port is part of the Fabric it is called an F_Port, otherwise it is an N_Port.

Fibre Channel is defined in terms of a hierarchy of functions or 'protocol layers'.

- FC-0: Physical Link
- FC-1: Transmission Protocol
- FC-2: Signaling and Framing Protocol
- FC-3: Common Services
- FC-4: Mapping

All layers except FC-4 are defined in

- “Fibre Channel Physical and Signaling Interface (FC-PH)” - ANSI

The Arbitrated Loop is an extension of FC-1 and FC-2 that allows more than two nodes to be connected without the expense of a Fabric. A port that connects to an Arbitrated Loop is generally referred to as an L_Port. However, the terms NL_Port and FL_Port are also used if it is necessary to distinguish whether or not the L_Port is part of a Fabric. FC-AL is defined in:

- “Fibre Channel Arbitrated Loop (FC-AL)” -ANSI

FC-4, the Mapping layer, defines how other communication protocols (e.g. SCSI, IPI-3, HIPPI) may use Fibre Channel functions. The FC-4 used by The Drive is SCSI-FCP, which is defined in:

- “Fibre Channel Protocol for SCSI (SCSI-FCP)” -ANSI

The 'drive Profile', recommends which features from the other documents should and should not be implemented in order to ensure compatibility between devices from different manufacturers. The full title of the drive Profile is:

- “Fibre Channel Private Loop SCSI Direct Attach (FC-PLDA)” -ANSI

An additional document, the ‘Public Loop Profile’, describes the additional features needed to support communication over a FC Fabric Topology. In addition to communicating with devices on their local loop, Public Loop devices can also communicate to devices across a network through the use of a “Fabric Port.” Public Loop behavior is defined in:

- “Fibre Channel - Fabric Loop Attach (FC-FLA)” -ANSI

The Drive complies with the FC-FLA, which makes the following requirements:

1. The upper two bytes of the device’s NL_Port ID must not be zero.
2. The device must be able to communicate with the Fabric Port i.e., open AL_PA 0x00.

In addition, the Drive implements FC-FLA as a super-set of the FC-PLDA. Thus, the Drive conforms to FC-PLDA except for those behaviors explicitly defined by the FC-FLA profile. The PLDA makes the following additional requirements.

1. The FC-4 is SCSI-FCP.
2. The Drive provides a Class 3 Fibre Channel service.
 - NO indication of (un)successful class 3 frame delivery is transmitted.
 - Frame flow control is buffer-to-buffer only.
 - Class 1 and 2 frames are ignored.
3. Direct point-to-point attachment with an N_Port is not supported.

15.1.1 Node and Port names

Every Fibre Channel Node and Port in the world must have a unique name. The Drive is a node with two L_Ports and therefore has a Node_Name and two Port_Names.

Both the Node_Name and Port_Name are in 64-bit IEEE Registered Name format, as illustrated in Table 20.

Table 20: IEEE Registered Name format

Bit				
63-60	59-36	35-24	23-2	1-0
0101	OUI in Canonical Form	Block Assignment	S/N	Object

The Name Address Authority field (0101b) specifies the format used for the rest of the name as follows:

- **OUI** Organizational Unique Identifier (24 bits). Canonical form means that each byte is stored in “bit reversed” order.
- **Block Assignment** Block assignment within HGST, a Western Digital Company
- **Object** Node/Port Identifier
 - 00b** Node
 - 01b** Port 1
 - 10b** Port 2
 - 11b** Not assigned
- **S/N** Sequentially increasing drive serial number assigned at manufacturing.

15.1.2 NL_Port address

An NL_Port address, as illustrated in Table 21, is a 3 byte value that uniquely identifies an NL_Port on a Fibre Channel network. The Domain and Area bytes are assigned to a NL_Port by its local FL_Port. The least significant byte in the NL_Port address is the Arbitrated Loop Physical Address (AL_PA), and uniquely identifies a NL_Port on its local loop.

A Hard AL_PA is one which is supplied, via the Drive connector, from external switches or jumpers. A Soft AL_PA is one which is assigned during Loop Initialization

Table 21: NL_Port address

Bits 23-16	Bits 15-8	Bits 7-0
Domain	Area	AL_PA

15.1.3 Primitive signals and sequences

The fundamental unit of transfer on a Fibre Channel link is the 8b/10b encoded Transmission Character. Only 256 characters are required to represent a byte of data so the set of valid transmission characters is sub-divided into Data Characters and Special Characters.

A Word is a group of four consecutive transmission characters.

An Ordered Set is a word that starts with a special character (to give word and byte sync) and which has special significance to the communication protocol. Fibre Channel defines the following types of ordered sets:

- Frame delimiters
- Primitive signals
- Primitive sequences

A Frame Delimiter immediately precedes or follows the contents of a frame (see 15.1.4, “Frames” on page 46). Fibre Channel defines the following:

- **SOF** Start of frame
- **EOF** End of frame

A Primitive Signal is an ordered set that has special meaning when received by itself. Fibre Channel defines the following:

- **R_RDY** Receiver Ready
- **IDLE** Idle

FC-AL adds the following:

- **ARBx** Arbitrate
- **ARB(F0)** Arbitrate
- **OPNyx** Open Full-duplex
- **OPNyy** Open Half-duplex
- **OPNfr** Open Broadcast Replicate¹
- **OPNyr** Open Selective Replicate¹
- **CLS** Close
- **MRKtx** Mark¹

A Primitive Sequence is an ordered set that is transmitted repeatedly and continuously. Three or more of the ordered sets must be received consecutively in order to recognize the primitive sequence. Fibre Channel defines the following:

- **NOS** Not Operational²
- **OLS** Offline²
- **LR** Link Reset²
- **LRR** Link Reset Received²

FC-AL adds the following:

- **LPByx** Loop Port Bypass
- **LPEyx** Loop Port Enable

1. Since normal buffer-to-buffer flow control is disabled when OPNr is used, there is no guarantee that the drive has a buffer available to receive an in-bound frame. It is therefore, recommended that OPNr not be used.

2. Used to convey information about a dedicated connection and therefore not relevant to FC-AL. If detected during an open connection, the drive will immediately close. Otherwise it will simply re-transmit.

- **LPEfx** Loop Port Enable all
- **LIP(F7,F7)** Loop Initialization, no valid AL_PA
- **LIP(F8,F7)** Loop Initialization, loop failure, no valid AL_PA
- **LIP(F7,AL_PS)** Loop Initialization, valid AL_PA
- **LIP(F8,AL_PS)** Loop Initialization, loop failure, valid AL_PA
- **LIP(AL_PD,AL_PS)** Loop Initialization, reset L_Port

15.1.4 Frames

Information transfer is achieved via frames that are constructed from words and ordered sets. All frames have the same general format, as shown in Table 22.

Table 22: General frame format

Field name	SOF	Header	Payload	Fill bytes	CRC	EOF
Field size, # of bytes	4	24	0 to 2048	0 to 3	4	4

The SOF ordered set indicates the start of frame and provides word synchronization.

The Header is the first field after the SOF delimiter. It is used by the link control facility to control link operations, control device protocol transfer, and to detect missing or out of order frames. The header is illustrated in Table 23.

FC-PH describes the content of each field except 'Parameter'. SCSI-FCP defines this field as RLTV_OFF (Relative Offset).

Table 23: Frame header

Byte word	3	2	1	0
0	R-CTL	D-ID		
1	Reserved	S-ID		
2	TYPE	F-CTL		
3	SEQ-ID	DF-CTL	SEQ-CNT	
4	OX-ID		RX-ID	
5	Parameter (RLTV_OFF)			

The Payload follows the header and has a length between 0 and 2048 bytes, which must be divisible by 4. An additional 0-3 fill bytes are appended to the payload in order to ensure that it ends on a word boundary.

The Cyclic Redundancy Check (CRC) is a four byte field following the payload. It is used to verify the integrity of the header and payload.

The EOF ordered set marks the end of a frame.

15.1.5 Sequences

A Sequence is a set of one or more related frames that flow in one direction only. The sequence is identified by the Sequence Identifier (SEQ_ID) field in the frame header.

15.1.6 Exchanges

An Exchange is a set of one or more related non-concurrent sequences that may flow in the same or opposite directions. The exchange is identified by an Originator Exchange Identifier (OX_ID) and a Responder Exchange Identifier (RX_ID) in the frame header.

15.2 Basic Link Services

The Basic Link Services are all frames with no payload. The Header TYPE field is set to 00h (Basic Link Service) and R_CTL is set to 1000xxxxb (Basic Link_Data, Code = xxxx).

Table 24: Basic link service command codes

Command	Description	Abbr.
0000	No Operation	NOP
0001	Abort Sequence	ABTS
0010	Remove Connection	RMC
0100	Basic Accept	BA_ACC
0101	Basic Reject	BA_RJT
Others	Reserved	

NOP and RMC are prohibited by the drive Profile. If the drive receives either an NOP or an RMC, it will ignore it.

Note: PLDA 2.1 specifies: “Reserved FC-PH fields are not required to be checked for zeroes. Validity bits set to 0 remove any requirement to check the corresponding field for zeroes (e.g., if F_CTL bit 3=0, receiving N_Ports are not required to verify that the parameter field in word 5 of the frame header contains zeroes).” As such, the drive does not validate 1) reserved FC fields or 2) fields that are not reserved but are not valid for the current frame (as the example above with F_CTL bit 3). This does not apply to any reserved field checking and testing within the FCP_CDB. These fields are checked as per ANSI SCSI requirements.

15.2.1 Abort sequence (ABTS)

Although ABTS is a Fibre Channel Basic Link Service, it is used by SCSI-FCP to implement the Abort Task, Task Management function. It may only be used by a SCSI initiator to abort an entire exchange using the Recovery Abort protocol. Refer to 16.4.1, “Abort Task (Implemented as ABTS BLS)” on page 101 for a description of the Recovery Abort Protocol and the frame payloads.

The response to ABTS is either BA_ACC or BA_RJT.

15.2.2 Basic accept (BA_ACC)

BA_ACC indicates that a Basic Link Service Request has been completed. The drive only sends a BA_ACC in response to an ABTS Basic Link Service.

Table 25: BA_ACC Payload

Byte	Item	Size (Bytes)
0	SEQ_ID Validity (80h = valid, 00h = invalid)	1
1	SEQ_ID	1
2	Reserved	2
3		
4	OX_ID	2
5		
6	RX_ID	2
7		
8	Low SEQ_CNT	2
9		
10	High SEQ_CNT	2
11		

- **SEQ_ID Validity** specifies whether the SEQ_ID field in the BA_ACC payload is valid or not. The drive always sets this field to 00h.
- **SEQ_ID** specifies the last SEQ_ID which is deliverable to the Upper Level Protocol. Since the drive uses the ABTS protocol to abort an entire exchange, this field is unused. The drive sets this field to 00h.
- **OXID** specifies the OXID of the exchange that has been aborted by the drive.
- **RXID** specifies the RXID of the exchange that has been aborted by the drive.
- **Low SEQ_CNT** specifies the last data frame of the last delivered sequence. Since the drive aborts the entire exchange, this field is set to 0000h.
- **High SEQ_CNT** is only valid when an ABTS is used to abort a single sequence. Since the drive only supports aborting of entire exchanges, this field is set to 0000h.

15.2.3 Basic reject (BA_RJT)

BA_RJT indicates that a Basic Link Service Request has been rejected. The payload contains a four byte reason code to indicate why the request was rejected.

Table 26: BA_RJT Payload

Byte	Item	Size (Bytes)
0	Reserved	1
1	Reason Code	1
2	Reason Explanation	1
3	Vendor Unique	1

Table 27: BA_RJT Reason Codes

Code	Description
0000 0001	Invalid command code
0000 0011	Logical error
0000 0101	Logical busy
0000 0111	Protocol error
0000 1001	Unable to perform command requested
1111 1111	Vendor unique error
Others	Reserved

Table 28: BA_RJT Reason Code Explanations

Code	Description
0000 0000	No additional explanation
0000 0011	Invalid OX_ID-RX_ID combination
0000 0101	Sequence aborted, no Sequence information provided
Others	Reserved

15.3 Extended Link Services

For Extended Link Service frames, the Header TYPE field is set to 01h (Extended Link Service). R_CTL is either 22h (Extended Link Data, Unsolicited Control) for a Request, or 23h (Extended Link Data, Solicited Control) for a Reply.

The first byte of the payload is the LS_Command and encodes the Request or Reply, as shown in Table 29 and Table 30.

Table 29: Extended Link Service replies

Code	Reply	Abbr.
02h	Accept	LS_ACC
01h	Link Service Reject	LS_RJT

Table 30 is a comprehensive list of all Extended Link Service commands supported by the drive.

Table 30: Extended Link Service requests

Code	Request	Abbr.
52h	Address Discovery	ADISC
60h	Fabric Address Notification	FAN
04h	Fabric Login	FLOGI
05h	Logout	LOGO
50h	Port Discovery	PDISC
03h	Port Login	PLOGI
20h	Process Login	PRLI
21h	Process Logout	PRLO
0Fh	Read Link Status	RLS
53h	Report Node Capabilities	RNC
78h	Report Node ID	RNID
12h	Re-instate Recovery Qualifier	RRQ
24h	Third Party Process Logout	TPRLO
7Dh	Report Port Speed Capabilities	RPSC
Note: Only lists ELSs supported by the drive. Refer to FC-PH for a complete list of ELSs.		

The code '11h' for the TEST ELS (not supported by the drive) is also used by the LIPxx frames that circulate during Loop Initialization. To differentiate these from a TEST ELS, the second byte of the payload must be examined and compared with Table 31.

Note: PLDA 1.10 specifies: “Reserved FC-PH fields are not required to be checked for zeroes. Validity bits set to 0 remove any requirement to check the corresponding field for zeroes (e.g., if F_CTL bit 3=0, receiving N_Ports are not required to verify that the parameter field in word 5 of the frame header contains zeroes).” As such, the drive does not validate 1) reserved FC fields or 2) fields that are not reserved but are not valid for the current frame (as shown in the example above with F_CTL bit 3). This does not apply to any reserved field checking and testing within the FCP_CDB. These fields are checked as per ANSI

SCSI requirements.

Table 31: Extended Link Service request 11h qualifiers

Code	Description	Abbr.
01h	Select Master	LISM
02h	Fabric Assign AL_PA	LIFA
03h	Previously Acquired AL_PA	LIPA
04h	Hard Assigned AL_PA	LIHA
05h	Soft Assigned AL_PA	LISA
06h	Report AL_PA Position Map	LIRP
07h	Loop AL_PA Position Map	LILP

If a Loop Initialization frame is received when the Port is not performing a Loop Initialization Procedure, it will be responded to by an LS_RJT containing a reason code of “Command not supported” and a reason code explanation of “No additional explanation”. This also applies to all unsupported link services, as defined by Table 30.

15.3.1 Link Service Accept (LS_ACC)

LS_ACC is used in response to an Extended Link Service Request. It indicates that the request has been completed.

The LS_ACC payload depends upon the Extended Link Service Request and is therefore described separately for each of the following:

- 15.3.3 “Port Login (PLOGI)”
- 15.3.4 “Logout (LOGO)”
- 15.3.5 “Fabric Login (FLOGI)”
- 15.3.6 “Fabric Address Notification (FAN)”
- 15.3.7 “Port Discovery (PDISC)”
- 15.3.8 “Address Discovery (ADISC)”
- 15.3.9 “Process Login (PRLI)”
- 15.3.10 “Process Logout (PRLO)”
- 15.3.11 “Read Link Error Status Block (RLS)”
- 15.3.12 “Report Node Capabilities (RNC)”
- 15.3.13 “Re-instate Recovery Qualifiers (RRQ)”
- 15.3.14 “Third Party Process Logout (TPRLO)”
- 15.3.15 “Request Node Identification Data (RNID)”
- 15.3.16 “Report Port Speed Capabilities (RPSC)”

15.3.2 Link Service Reject (LS_RJT)

LS_RJT indicates that the Extended Link Service request has been rejected. The payload (shown in Table 32) contains a Reason Code and a Reason Explanation.

Table 32: LS_RJT payload

Byte	Item	Size (Bytes)
0 3	0100 0000h	4
4	Reserved	1
5	Reason Code (Table 33)	1
6	Reason Explanation (Table 34)	1
7	Vendor Unique	1

Table 33: LS_RJT reason codes

Code	Description
01h	Invalid command code
03h	Logical error
05h	Logical busy
07h	Protocol error
09h	Unable to perform command requested
0Bh	Command not supported
FFh	Vendor unique error

Table 34: LS_RJT reason code explanations

Code	Description
00h	No additional explanation
01h	Service Parm error - Options
03h	Service Parm error - Initiator Ctl
05h	Service Parm error - Recipient Ctl
07h	Service Parm error - Rec Data Field Size
09h	Service Parm error - Concurrent Seq
0Bh	Service Parm error - Credit
0Dh	Invalid Port Name
0Eh	Invalid Node/Fabric Name
0Fh	Invalid Common Service Parameters
19h	Command (request) already in progress
1Fh	Invalid N_Port identifier
29h	Insufficient resources to support Login
2Ah	Unable to supply requested data
2Ch	Request not supported

Note: Refer to FC-PH for a full list. Only the explanations relevant to supported Extended Link Services are shown.

15.3.3 Port Login (PLOGI)

PLOGI is used by the Initiator to register Service Parameters with the Target if Implicit Login is not being used (see 17.10.11, “Mode Page 19h (Fibre Channel Port Control Page)” on page 172). The Target responds with an LS_ACC that has a payload similar to the PLOGI but which contains the Targets parameters. This exchange of parameters establishes the operating environment between the Initiator and the Target.

PLOGI can also be issued by the Target when it is operating as a Public Loop Device. In this case, the Target uses a PLOGI to establish the operating environment between it and the Fabric Name Server.

Bytes 1, 2, and 3, of the payload must be set to zeros. Otherwise, the drive will respond with an LS_RJT containing a reason code of “Command not supported”, and a reason code explanation of 00h “No additional explanation”.

Table 35: PLOGI_REQ/PLOGI_ACC payload

Byte	Item	Size (Bytes)
0-3	Request = 0300 0000h; Accept = 0200 0000h	4
4 - 19	Common Service Parameters	16
20 - 27	Port Name	8
28 - 35	Node Name	8
36 - 51	Class 1 Service Parameters	16
52 - 67	Class 2 Service Parameters	16
68 - 83	Class 3 Service Parameters	16
84 - 99	Reserved	16
100 - 115	Vendor Version Level	16

15.3.3.1 Common Service Parameters

The Common Service Parameters apply to all classes of service and are exchanged during Login. The table below defines the applicability, by class as well as by PLOGI, FLOGI, PLOGI LS_ACC and FLOGI LS_ACC, of the Common Service Parameters to N_Port and Fabric Login. These are words 1-4 in the Payload.

Table 36: Common Service Parameter applicability (part 1 of 2)

Service parameter	Word	Bits	PLOGI and PLOGI LS_ACC Parameter applicability				FLOGI Parameter applicability				FLOGI LS_ACC Parameter applicability			
			Class				Class				Class			
			1*	2	3	4	1*	2	3	4	1*	2	3	4
FC-PH Version - Obsolete	1	31-16	n	n	n	n	n	n	n	n	n	n	n	n
Buffer-to-Buffer Credit	1	15-0	y	y	y	n	y	y	y	n	y	y	y	n
Common Features	2	31-16												
Continuously increasing relative offset	2	31	y	y	y	y	n	n	n	n	n	n	n	n
Clean Address	2	31	n	n	n	n	n	n	n	n	y	y	y	y
Random relative offset	2	30	y	y	y	y	n	n	n	n	n	n	n	n
Valid Vendor Version Level	2	29	y	y	y	y	y	y	y	y	n	n	n	n
N_Port/F_Port	2	28	y	y	y	y	y	y	y	y	y	y	y	y
BB_Credit Management	2	27	y	y	y	y	y	y	y	y	n	n	n	n
E_D_TOV Resolution	2	26	y **	y **	y **	y **	n	n	n	n	y	y	y	y
Multicast supported by Fabric	2	25	n	n	n	n	n	n	n	n	y	y	y	y
Broadcast supported by Fabric	2	24	n	n	n	n	n	n	n	n	y	y	y	y
Hunt Group routing supported by Fabric	2	23	n	n	n	n	n	n	n	n	y	y	y	y
Simplex Dedicated Connection - Obsolete	2	22	n	n	n	n	n	n	n	n	n	n	n	n
Reserved for security	2	21	n	n	n	n	n	n	n	n	n	n	n	n
Clock Synchronization Primitive Capable	2	20	y	y	y	y	y	y	y	y	y	y	y	y
R_T_TOV Value	2	19	y	y	y	y	y	y	y	y	y	y	y	y
Dynamic Half Duplex Supported	2	18	y	y	y	y	y	y	y	y	y	y	y	y
SEQ_CNT	2	17	y	y	y	y	n	n	n	n	n	n	n	n
Payload bit	2	16	y	y	y	y	y	y	y	y	y	y	y	y

"y" indicates yes, applicable (i.e. has meaning)

"n" indicates no, not applicable (i.e. has no meaning)

* The Class 1 Service Parameters shall be used for Class 6. Each has the same applicability as Class 1.

** E_D_TOV resolution and the corresponding value are only meaningful in a point-to-point topology and when doing PLOGI with an NL_Port on the same loop.

Table 37: Common Service Parameter applicability (part 2 of 2)

Service parameter	Word	Bits	PLOGI and PLOGI LS_ACC Parameter applicability				FLOGI Parameter applicability				FLOGI LS_ACC Parameter applicability			
			Class				Class				Class			
			1*	2	3	4	1*	2	3	4	1*	2	3	4
BB_SC_N	2	15-12	y	y	y	n	y	y	y	n	y	y	y	n
Buffer-to-Buffer Receive Data Field Size	2	11-0	y	y	y	y	y	y	y	y	y	y	y	y
Nx_Port Total Concurrent Sequences	3	31-16	y	y	y	y	n	n	n	n	n	n	n	n
Relative offset by Info Category	3	15-0	y	y	y	y	n	n	n	n	n	n	n	y
R_A_TOV	3	31-0	n	n	n	n	n	n	n	n	y	y	y	y
E_D_TOV Value	4	31-0	y**	y**	y**	y**	n	n	n	n	y	y	y	y

"y" indicates yes, applicable (i.e. has meaning)

"n" indicates no, not applicable (i.e. has no meaning)

* The Class 1 Service Parameters shall be used for Class 6. Each has the same applicability as Class 1.

** E_D_TOV resolution and the corresponding value are only meaningful in a point-to-point topology and when doing PLOGI with an NL_Port on the same loop.

The upper byte of the FC-PH Version field indicates the highest version of FC-PH that is supported and the lower byte indicates the lowest. The code/version relationship is shown in Table 38.

Table 38: FC-PH Version

Hex value	Version
00	None
06	FC-PH 4.0
07	FC-PH 4.1
08	FC-PH 4.2
09	FC-PH 4.3
10h	FC-PH-2
20h	FC-PH-3
Others	Reserved

The **BB_Credit** field indicates the number of frame buffers that a port guarantees to have immediately available when a loop circuit is opened. The drive returns 0 (zero) in this field; i.e., every Initiator is given a Login BB_Credit of 0.

The **Common Features** field contains flags.

- **CIO - Continuously Increasing Offset**

When set to one, this flag indicates that the Port supports Continuously Increasing Relative Offset within a Sequence on a frame by frame SEQ_CNT basis. The Relative Offset is only present if bit 3 of the F_CTL field in the frame header is set to 1b. The drive requires that an initiator support Continuously Increasing Relative Offset. If this field is

not one, the drive will respond with an LS_RJT containing a reason code of 03h “Logical error”, and a reason code explanation of 0Fh “Invalid Common Service Parameters”. The drive returns a 1b in this field.

- **RRO - Random Relative Offset**

When set to one, this flag indicates that the Port supports Random Relative Offset within a Sequence. The Relative Offset is only present if bit 3 of the F_CTL field in the frame header is set to 1b. The drive returns a 0b in this field.

- **VV - Valid Vendor Version Level**

When set to one, this flag indicates that the Vendor Version Level field in the frame payload contains valid information. The drive returns a 0b in this field.

- **N/F Port - N_Port/F_Port**

When set to one, this flag indicates that the Port is an F_Port. When set to zero, it is an N_Port. The drive returns a 0b in this field.

- **ABCM - Alternate BB-Credit Model**

When set to one, this flag indicates that the Port supports the Alternate BB-Credit model. The drive returns a 1b in this field.

- **E_D_TOV - E_D_TOV Resolution**

When set to one, this flag indicates that the resolution of the E_D_TOV timer shall be 1 ns. When set to zero, the resolution shall be 1 ms. The flag is only applicable in a point to point topology. The drive returns a 0b in this field.

- **DS - Dedicated Simplex**

Only applies to class 1 service. The drive sets this field to 0b.

- **DHD - Dynamic Half Duplex**

When set to one, this flag indicates that the Port supports Dynamic Half Duplex. The drive returns a 0b in this field.

- **SEQ_CNT - SEQ_CNT**

When set to one, this flag indicates that the Port is guaranteeing that it will transmit all frames within an Exchange using a continuously increasing SEQ_CNT. If set to zero, normal FC_PH rules regarding SEQ_CNT usage apply. The drive returns a 0b in this field.

- **PL - Payload Length**

When set to one, this flag indicates that PLOGI Payload Length shall be 256 bytes. If set to zero, the flag specifies that the normal 116 byte PLOGI Payload specified in FC-PH-2 is used. The drive returns a 0b in this field.

The **Buffer-to-Buffer Receive Data_Field Size** field (word 2, bits 11-0) specifies the largest FT_1 frame Data_Field Size that may be received by the Nx_Port supplying the Service Parameters as a Sequence Recipient for:

- a) a connect-request (SOFc1),
- b) a Class 2 Data frame, or
- c) a Class 3 Data frame

Values less than **256** or greater than 2^{112} are invalid and shall be a multiple of four bytes. An Fx_Port shall support a Data Field size of at least 256 bytes.

The **Total Concurrent Sequences** field indicates the number of concurrent sequences that the Port can support for all 3 classes. The value sent must be > 0 , or the drive will respond with an LS_RJT with a reason code of 03h “Logical error” and a reason code explanation of 09h “Service Parm error - Concurrent Seq”. The drive returns a value of FFh in this field.

The **Relative Offset by Information Category** field is all flags. The bit position of each set (1) bit determines the Information Category; e.g., if Relative Offset is supported for category 0001b (Solicited Data), bit 1 is set. The drive returns a value of 0002h in this field (Relative Offset supported for Solicited Data only).

The **E_D_TOV** field specifies the **E_D_TOV** (Error Detect Timeout Value) in units of 1ns or 1ms. If the **E_D_TOV** Resolution bit is set to one, this field specifies **E_D_TOV** as a count of 1ns increments, otherwise it specifies **E_D_TOV** as a count of 1ms increments. It is only applicable in a point to point topology. The drive returns zero in this field. It is not interpreted or checked.

15.3.3.2 Class Service Parameters

Since the drive only supports Class 3 service, it only checks and responds to the Class 3 Service Parameters.

Table 39: Class Service Parameters

Byte	BIT								Default ACC Values
	7	6	5	4	3	2	1	0	
68	CV	IM	TM	LDM	SD	DS	Camp-On	BufC1	80h
69	Priority	Reserved							00h
Initiator Control Flags (Bytes 70-71)									
70	X_ID Reassign		IP Associator		ACK_0	ACK_N	ACK Gen	DCC	00h
71	DCHBS		DEC	CSC	Reserved				00h
Recipient Control Flags (Bytes 72-73)									
72	ACK_0 Cap	ACK_N Cap	X_ID Int.	EP Supported		Reserved	Categ. Per Sequence		00h
73	DCC	DCHBS		DDC	CSS	Reserved			00h
74	(MSB) Receive Data-Field Size								08h
75	(LSB)								00h
76	(MSB) Concurrent Sequences								00h
77	(LSB)								FFh
78	(MSB) End-to-End Credit								00h
79	(LSB)								00h
80	(MSB) Open Sequences per Exchange								00h
81	(LSB)								01h
82	(MSB) Class 6 Multi-cast RX_ID								00h
83	(LSB)								00h

The **Service Options** field contains the following flags.

- **CV - Class Validity**

When set to one, this flag indicates that the class of service is supported. The drive only supports Class 3, thus this bit is only set for the Class 3 Service Parameters. If this bit is not set for Class 3 Service Parameters, the drive will respond with an LS_RJT containing a reason code of 03h “Logical error” and a reason code explanation of 01h “Service Parm error - Options”.

- **IM - Intermix Mode**

The flag only applies to class 1 service and is reserved for classes 2 and 3. The drive returns 0b in this field.

- **TM - Transparent Mode**

The flag does not apply to PLOGI. The drive returns 0b in this field.

- **LDM - Lock Down Mode**

This flag does not apply to PLOGI. The drive returns 0b in this field.

- **SD - Sequential Delivery**

This flag does not apply to PLOGI. The drive returns 0b in this field.

- **DS - Dedicated Simplex**

This flag only applies to class 1 service. The drive returns 0b in this field.

- **Camp-On - Camp-On**

This flag only applies to class 1 service. The drive returns 0b in this field.

- **BufC1 - Buffered Class 1**

This flag only applies to class 1 service. The drive returns 0b in this field.

- **Priority**

This flag does not apply to class 3 service. The drive returns 0b in this field.

The **Initiator Control** flags specify which protocols, policies, or functions the supplier of the Service Parameters requests of the recipient or is capable of.

- **X_ID Reassign - X_ID Reassignment**

- 0 0** X_ID Reassignment not supported.
- 0 1** X_ID Reassignment supported.
- 1 0** Reserved.
- 1 1** X_ID Reassignment required.

The X_ID Reassignment bits only apply to class 1 or 2 and therefore are not checked by the drive.

- **IP Associator - Initial Process Associator**

- 0 0** Initial Process Associator not supported.
- 0 1** Initial Process Associator supported.
- 1 0** Reserved.
- 1 1** Initial Process Associator required.

The drive accepts values of “00” and “01” in this field. If other values are sent, the drive responds with an LS_RJT containing a reason code of 03h “Logical error” and a reason code explanation of 03h “Service Parm error - Initiator Ctl”. The drive returns 00b in this field.

- **ACK_0 - ACK_0 Capability**

This flag does not apply to class 3 and therefore is not checked by the drive. The drive returns 0b in this field.

- **ACK_N - ACK_N Capability**

This flag does not apply to class 3 and therefore is not checked by the drive. The drive returns 0b in this field.

- **ACK Gen - ACK generation assistance**

This flag does not apply to class 3 and therefore is not checked by the drive. The drive returns 0b in this field.

- **DCC - Data compression capable**

When set to one, this flag indicates the Port supports data compression as a Sequence Initiator. The drive does not support data compression, therefore the drive returns 0b in this field.

- **DCHBS - Data compression History buffer size**

This field indicates the History buffer size supported by the Port as a Sequence Initiator. The drive does not support data compression, therefore, the drive returns 00b in this field.

- **DEC - Data encryption capable**

When set to one, this flag indicates the Port supports data encryption as a Sequence Initiator. The drive does not support data compression, therefore the drive returns 0b in this field.

- **CSC - Clock synchronization capable**

When set to one, this flag indicates the Port is capable of performing clock synchronization as a Sequence Initiator (Clock Synchronization Server). The drive is not capable of performing clock synchronization, therefore the drive returns 0b in this field.

The **Recipient Control** flags specify which functions the supplier of the Service Parameters supports when receiving frames.

- **ACK_0 Cap - ACK_0 Capability**

This flag only applies to class 1 or 2, therefore is not checked by the drive. The drive returns 0b in this field.

- **ACK_N Cap - ACK_N Capability**

This flag only applies to class 1 or 2, therefore is not checked by the drive. The drive returns 0b in this field.

- **X_ID Int. - X_ID Interlock**

This flag only applies to class 1 or 2, therefore is not checked by the drive. The drive returns 0b in this field.

- **EPS - Error Policy Supported**

0 0	Only discard supported
0 1	Reserved
1 0	Both discard and process supported
1 1	Reserved

The Error Policy bits are not checked by the drive. The drive returns 00b in this field.

- **Categ. per Sequence - Categories per Sequence**

0 0	1 Category/Sequence
0 1	2 Categories/Sequence
1 0	Reserved
1 1	More than 2 Categories/Sequence. The drive returns 00b in this field.

The **Receive Data Field** size is the largest payload (in bytes) that the Port is capable of receiving. It must be less than or equal to the Buffer to Buffer Receive Data Field specified in the Common Service Parameters. Sizes less than 128, greater than 2112, or not divisible by 4 are invalid, in which case the drive will respond with an LS_RJT containing a reason code of 03h “Logical error” and a reason code explanation of 07h “Service Parm error - Rec Data Field Size”. The drive returns 2048 in this field.

The **Concurrent Sequences** field specifies the number of separate Sequences that the drive is capable of tracking. The value sent must be > 0, or the drive will respond with an LS_RJT containing a reason code of 03h “Logical error” and a reason code explanation of 09h “Service Parm error - Concurrent Seq”. The drive returns a value of 255 in this field.

The **End-to-end Credit** field applies only to classes 1 and 2 and is therefore not checked by the drive. The drive returns 0 in this field.

The **Open Sequences per Exchange** field indicates the maximum number of Sequences per Exchange that can be open at the recipient at one time. The value sent must be > 0, or the drive will respond with an LS_RJT containing reason code of 03h “Logical error” and a reason code explanation of 00h “No additional explanation”. The drive returns 01b in this field.

The **Class 6 Multicast RXID** is used in Class 6 only and is therefore not checked by the drive. The drive returns 00b in this field.

15.3.4 Logout (LOGO)

LOGO is used by an Initiator to request invalidation of the Service Parameters that were exchanged during PLOGI. If implicit login is enabled by the settings in Mode Page 19, then the default Service Parameters apply. If no valid Service Parameters exist for an Initiator, a LOGO_ACC is still returned.

Bytes 1, 2, and 3 of the payload must be set to zero. Otherwise, the drive will respond with an LS_RJT containing a reason code of 0Bh “Command not supported” and a reason code explanation of 00h “No additional explanation”.

Table 40: LOGO payload

Byte	Item	Size (Bytes)
0-3	0500 0000h	4
4	Reserved	1
5-7	Port Identifier	3
8 - 15	Port Name	8

Table 41: LOGO ACC payload

Byte	Item	Size (Bytes)
0-3	0200 0000h	4

15.3.5 Fabric Login (FLOGI)

FLOGI is used by the drive to register its Service Parameters with the Fabric assuming the drive exists on a Public Loop (i.e. the loop contains an FL_Port). Once the drive has completed FLOGI, it is then permitted to use the local FL_Port as a gateway, allowing it to communicate with other N_Ports and NL_Ports attached to the fabric. (See 19.5, “Public Loop Operation” on page 276.)

The drive performs a FLOGI by opening a full duplex connection to FL_Port located at AL_PA 00h and sending a FLOGI request to the well-known address of FFFFFFFh. The FL-Port responds to a valid request with a FLOGI Accept frame that contains the FL_Port’s operating parameters. The D-ID field of the FLOGI Accept frame contains the fabric assigned Domain, Area, and the AL-PA of the drive performing the FLOGI. The drive uses this address in all further communication with other N_Ports and Public NL_Ports. The drive will not perform a FLOGI if the drive fails to detect an FL_Port during Loop-Initialization.

Table 42: FLOGI_REQ/FLOGI_ACC payload

Byte	Item	Size (Bytes)
0-3	Request = 0400 0000h; Accept = 0200 0000h	4
4 - 19	Common Service Parameters	16
20 - 27	Port Name	8
28 - 35	Node Name	8
36 - 51	Class 1 Service Parameters	16
52 - 67	Class 2 Service Parameters	16
68 - 83	Class 3 Service Parameters	16
84 - 99	Reserved	16
100 - 115	Vendor Version Level	16

15.3.5.1 Common Service Parameters

The Common Service Parameters apply to all classes of service and are exchanged during Login.

Table 43: Common Service Parameters (FLOGI_REQ/FLOGI_ACC)

Byte	BIT								Default Request Values
	7	6	5	4	3	2	1	0	
4	PH Version - Highest Supported								20h
5	PH Version - Lowest Supported								20h
6	(MSB) Buffer-to-Buffer Credit (LSB)								00h
7									00h
8	Reserved	VV	N/F Port	ABCM	Reserved			00h	
9	Reserved				DHD	Reserved	PL	00h	
10	(MSB) BB Received Data Field Size (LSB)								08h
11									00h
12	(MSB) FLOGI_REQ = Reserved; FLOGI_ACC = R_A_TOV (LSB)								00h
13									00h
14									00h
15									00h
16	(MSB) FLOGI_REQ = Reserved; FLOGI_ACC = E_D_TOV (LSB)								00h
17									00h
18									00h
19									00h

The upper byte of the **FC-PH Version** field indicates the highest version of FC-PH that is supported and the lower byte indicates the lowest. The code/version relationship is given by Table 38.

The **BB_Credit** field indicates the number of frame buffers that a port guarantees to have immediately available when a loop circuit is opened. The drive sets this field to 0 (zero); i.e., the drive grants the FL_Port a Login BB_Credit of 0.

The **Common Features** field contains the following flags:

- **VV- Valid Vendor Version Level**

When set to one, this flag indicates that the Vendor Version Level field in the frame payload contains valid information. The drive sets this field to 0b.

- **N/F Port - N_Port/F_Port**

When set to one, this flag indicates that the Port is an F_Port. When set to zero, it is an N_Port. The drive checks this field in the FLOGI_ACC and will revert to PLDA behavior if it is not set to one. The drive sets this field to 0b.

- **ABCM - Alternate BB-Credit model**

When set to one, this flag indicates that the Port supports the Alternate BB-Credit model. As per the FC-AL Specification, an L_Port and must support the Alternate BB-Credit model. The drive checks the value returned by the FL_Port and will revert to PLDA behavior if it is not set to one. The drive sets this field to 1b.

- **DHD - Dynamic Half Duplex**

When set to one, this flag indicates that the Port supports the Dynamic Half Duplex. The drive sets this field to 0b.

- **PL - Payload Length**

When set to one, this flag indicates that FLOGI Payload Length shall be 256 bytes. If set to zero, the flag specifies that the normal 116 byte FLOGI payload specified in FC-PH-2. The drive sets this field to 0b.

The **Buffer to Buffer Receive Data Field** size indicates the largest frame payload (in bytes) that the Port can receive. Sizes less than 128, greater than 2112, or not divisible by 4 are invalid. The drive check this field in the FLOGI_ACC payload and will revert to PLDA behavior if an invalid value is returned. The drive sets this field to 2048.

The **R_A_TOV** field specifies the R_A_TOV (Resource Allocation Timeout Value) supplied by the fabric in units of 1ms. The fabric port sets this field in its FLOGI_ACC payload. This field is reserved in the FLOGI_REQ and the drive therefore sets this field to zero. After FLOGI is complete, the drive will use this value for R_A_TOV until the drive is reset, reverts to PLDA behavior, or performs another FLOGI.

The **E_D_TOV** field specifies the E_D_TOV (Error Detect Timeout Value) supplied by the fabric in units of 1ns or 1ms. The fabric port sets this field in its FLOGI_ACC payload. This field is reserved in the FLOGI_REQ and the drive therefore sets this field to zero. After FLOGI is complete, the drive will use this value for E_D_TOV until the drive is reset, reverts to PLDA behavior, or performs another FLOGI.

15.3.5.2 Class Service Parameters

Since the drive only supports class 3 service, it only checks and responds to the class 3 Service Parameters.

Table 44: Class Service Parameters

Byte	BIT								Default Request Values
	7	6	5	4	3	2	1	0	
68	CV	IM	TM	LDM	SD	DS	Camp-On	BufC1	80h
69	Priority	Reserved							00h
Initiator Control Flags (Bytes 70-71)									
70	Reserved								00h
71	Reserved								00h
Recipient Control Flags (Bytes 72-73)									
72	Reserved								00h
73	Reserved								00h
74	Receive Data-Field Size (Reserved)								00h
75									(LSB)
76	Concurrent Sequences (Reserved)								00h
77									(LSB)
78	End-to-End Credit (Reserved)								00h
79									(LSB)
80	Open Sequences per Exchange (Reserved)								00h
81									(LSB)
82	Class 6 Multi-cast RXID (Reserved)								00h
83									(LSB)

The **Service Options** field contains the following flags:

- **CV - Class Validity**

When set to one, this flag indicates that the class of service is supported. The drive only supports class 3 thus this bit is only set for the class 3 Service Parameters. If the FLOGI_ACC indicates that the fabric does not support class 3, the drive will revert to PLDA behavior.

- **IM - Intermix Mode**

This flag only applies to class 1 service. It is reserved for classes 2 and 3. The drive sets this field to 0b.

- **TM - Transparent Mode**

This flag only applies to class 1 service. It is reserved for classes 2 and 3. The drive sets this field to 0b.

- **LDM - Lock Down Mode**

This flag only applies to class 1 service. It is reserved for classes 2 and 3. The drive sets this field to 0b.

- **SD - Sequential Delivery**

The drive sets this field to 1b to request that the fabric deliver all frames in the same order they were transmitted. The fabric sets this flag to 1b if it can honor this request. Since the drive does not support out of order frame delivery, the drive checks the FLOGI_ACC to ensure that this field is set to 1b and will revert to PLDA behavior if it is not.

- **DS - Dedicated Simplex**

This flag only applies to class 1 service. The drive sets this field to 0b.

- **Camp On - Camp-On**

This flag only applies to class 1 service. The drive sets this field to 0b.

- **BufC1 - Buffered Class 1**

This flag only applies to class 1 service. The drive sets this field to 0b.

- **Priority - Priority**

This flag does not apply to class 3 service. The drive sets this field to 0b.

The **Initiator Control Flags** are not meaningful for FLOGI and are therefore reserved. The drive sets this field to 00h.

The **Recipient Control Flags** are not meaningful for FLOGI and are therefore reserved. The drive sets this field to 00h.

The **Receive Data Field Size** field is not meaningful for FLOGI and is therefore reserved. The drive sets this field to 00h.

The **Concurrent Sequences** field is not meaningful for FLOGI and is therefore reserved. The drive sets this field to 00h.

The **End-to-End Credit** field is not meaningful for FLOGI and is therefore reserved. The drive sets this field to 00h.

The **Open Sequences per Exchange** field is not meaningful for FLOGI and is therefore reserved. The drive sets this field to 00h.

The **Class 6 Multi-cast RXID** field is not meaningful for FLOGI and is therefore reserved. The drive sets this field to 00h.

15.3.6 Fabric Address Notification (FAN)

The FAN ELS is sent by an FL_Port to all previously logged in (FLOGI) NL_Ports after an initialization event has occurred. An initialization event can be a LIP or any other event that may cause a Port to change its ID. The FAN ELS is used by attached NL_Ports to validate their operating parameters with the local FL_Port after an initialization event has occurred. The drive does not return a response to this ELS.

(Note: When operating as a private loop device, the drive will discard the FAN ELS request.)

Table 45: FAN payload

Byte	Item	Size (Bytes)
0-3	Command Code = 60000000h	4
4	reserved	1
5-7	Loop Fabric Address	3
8 -15	Fabric Port Name	8
16 -23	Fabric Name	8

The **Command Code** must be equal to 60000000h or the drive will ignore the request.

The **Loop Fabric Address** is the 3 byte Port_ID of the local FL_Port.

The **Fabric Port Name** is the world-wide-unique 8-byte Port_Name of the local FL_Port.

The **Fabric Name** is the world-wide-unique 8-byte name of the Fabric.

15.3.7 Port Discovery (PDISC)

An Initiator uses PDISC to exchange service parameters without affecting the operating environment between it and the drive. The PDISC and corresponding ACC payloads are exactly as defined for PLOGI (see 15.3.3, “Port Login (PLOGI)” on page 53), except that byte 0 of the payload is 50h. It can be used as a means of authentication following a Loop Initialization process. If AL_PA's of the Initiator and Target have not changed since the previous login, the Initiator and Target can continue where they left off prior to the LIP and all open exchanges will continue. If the current AL_PA's do not match the login values, the drive will implicitly log out the Initiator.

If some means of authentication following a LIP does not occur within RR_TOV, the drive will implicitly log out the Initiator.

Bytes 1, 2, and 3 of the payload must be set to zero. Otherwise, the drive will respond with an LS_RJT containing a reason code of 0Bh “Command not supported” and a reason code explanation of 00h “No additional explanation”.

15.3.8 Address Discovery (ADISC)

The ADISC ELS allows communicating N-Ports to exchange addresses and port/node name identifiers. It can be used as a quick means of authentication following a Loop Initialization process. If the AL_PA's of the Initiator and Target have not changed since the previous login, the Initiator and Target can continue where they left off prior to the LIP, and all open exchanges will continue. If the current AL_PA's do not match the login values, the drive will implicitly log out the Initiator. ADISC can also be used as a means to determine if a port was able to acquire its hard AL_PA during LIP.

Bytes 1, 2, and 3 of the payload must be set to zero. Otherwise, the drive will respond with an LS_RJT containing a reason code of 0Bh “Command not supported” and a reason code explanation of 00h “No additional explanation”.

If some means of authentication following a LIP does not occur within RR_TOV, the drive will implicitly log out the Initiator.

The following tables show the ADISC payload and the ADISC ACC payload.

Table 46: ADISC payload

Byte	Item	Size (Bytes)
0-3	Command Code = 52000000h	4
4	Reserved	1
5 -7	Hard Address of Originator	3
8 -15	Port Name of Originator	8
16 -23	Node Name of Originator	8
24	Reserved	1
25-27	N_Port ID of Originator	3

Table 47: ADISC ACC payload

Byte	Item	Size (Bytes)
0-3	ACC Code = 02000000h	4
4	Reserved	1
5-7	Hard Address of Responder	3
8 -15	Port Name of Responder	8
16 -23	Node Name of Responder	8
24	Reserved	1
25-27	N_Port ID of Responder	3

Hard Address: This 3 byte identifier consists of:

- The MSB is an 8-bit domain address. If the drive is acting as a Public Loop Device, it returns the domain address given to it by the Fabric port during login. Otherwise the drive returns 0s in this field.
- The middle byte is an 8-bit area address. If the drive is acting as a Public Loop Device, it returns the area address given to it by the Fabric port during login. Otherwise the drive returns 0s in this field.
- The LSB is the 8-bit AL_PA that the port attempts to acquire during the LIHA sequence of LIP. For the drive, this number is calculated from the SEL-ID pins on the SCA-2 backplane connector.

When the Hard Address field is equal to the N_Port ID field, the port was able to obtain its hard address during LIP. The drive does not check the value sent from the originator.

Port Name: IEEE unique address assigned during the manufacturing process.

Node Name: IEEE unique address assigned during the manufacturing process.

N_Port ID: This is the 24-bit NL_Port Identifier used in the S_ID of the ADISC Accept header. The lower 8 bits are the AL_PA the drive acquired during loop initialization, and the upper 16 bits are the domain and area addresses the drive obtained from the Fabric port. For private devices the upper two bytes should be all 0s. When this field matches the Hard Address field, the drive was able to acquire its hard AL_PA during LIP.

15.3.9 Process Login (PRLI)

The PRLI request informs the recipient of the capabilities and requirements of the originator. The recipient responds with an ACC to indicate agreement or LS_RJT otherwise.

Table 48: PRLI payload

Byte	Item	Size (Bytes)
0	20h	1
1	Page Length = 10h	1
2-3	Payload Length	2
4-max	Login service parameter pages	n*16

Table 49: PRLI ACC payload

Byte	Item	Size (Bytes)
0	02h	1
1	Page Length=10h	1
2-3	Payload Length	2
4-max	Login response service parameter pages	n*16

15.3.9.1 Process Login Service Parameter page

Table 50: Login Service Parameter page

Byte	Item	Size (Bytes)
0	TYPE Code = 08h	1
1	TYPE Code Extension = 00h	1
2-3	Flags	2
4-7	Originator Process Associator	4
8-11	Responder Process Associator	4
12-15	Service Parameters	4

The **TYPE code** field for SCSI-FCP is 08h, and the **TYPE Code Extension** field is 00h.

The **Flags** field is as follows:

- **Bit 15 - Originator Process Associator Valid**

When set to one, the Originator Process Associator field for this Service Parameter page is valid.

- **Bit 14 - Responder Process Associator Valid**

When set to one, the Responder Process Associator field for this Service Parameter page is valid.

- **Bit 13 - Establish Image Pair**

When set to one, this flag indicates that the Originator wishes to establish an image pair.

- **Bits 12-0 - Reserved**

The **Originator Process Associator** identifies a group of related processes (an 'image') within the originator.

The **Responder Process Associator** identifies a group of related processes (an 'image') within the responder.

The **Service Parameters** field contains flags, as follows:

- **Bits 31-7 - Reserved**

- **Bit 6 - Data Overlay Allowed**

When set to one, this flag indicates that the initiator function is capable of supporting data overlay.

- **Bit 5 - Initiator Function**

When set to one, this flag indicates that the process defined by this page is operating as a SCSI Initiator.

This bit must be set to one in order for the drive to accept the login request. If not, the drive will respond with an LS_RJT containing a reason code of 03h “Logical error” and a reason code explanation of 00h “No additional explanation”.

- **Bit 4 - Target Function**

When set to one, this flag indicates that the process defined by this page is operating as a SCSI Target.

- **Bit 3 - Command/Data Mixed Allowed**

When set to one, this flag indicates that FCP_CMND and FCP_DATA may be combined in one IU.

- **Bit 2 - Data/Response Mixed Allowed**

When set to one, this flag indicates that FCP_DATA and FCP_RSP may be combined in one IU.

- **Bit 1 - Read XFER_RDY Disabled**

When set to one, this flag indicates that the FCP_XFER_RDY IU may not be used for SCSI READ operations.

- **Bit 0 - Write XFER_RDY Disabled**

When set to one, this flag indicates that the FCP_XFER_RDY IU may not be used for SCSI WRITE operations.

15.3.9.2 Process Login Response Service Parameter page

Table 51: Login Response Service Parameter page

Byte	Item	Size (Bytes)	Drive Response
0	TYPE Code	1	08h
1	TYPE Code Extension	1	00h
2-3	Flags	2	21 00h
4-7	Originator Process Associator	4	00 00 00 00h
8-11	Responder Process Associator	4	00 00 00 00h
12-15	Service Parameters	4	00 00 00 12h

The **TYPE Code** field for SCSI-FCP is 08h, and the **TYPE Code Extension** field is 00h.

The **Flags** field is as follows:

- **Bit 15 - Originator Process Associator Valid**

Since the drive does not support Process Associators, this bit must be set to zero in the payload. If not, the drive will respond with a PRLI ACC response code of 07h. The drive returns 0b in this field.

- **Bit 14 - Responder Process Associator Valid**

Since the drive does not support Process Associators, this bit must be set to zero in the payload. If not, the drive will respond with a PRLI ACC response code of 04h. The drive returns 0b in this field.

- **Bit 13 - Image Pair Established**

The drive returns a 1b in this field.

- **Bit 12 - Reserved**

- **Bits 11-8 - Accept Response Code**

See Table 52.

- **Bits 7-0 - Reserved**

The **Originator Process Associator** and **Responder Process Associator** fields are not used and will be set to 00000000h by the drive.

The **Service Parameters** field contains flags, as follows:

- **Bits 31-6 - Reserved**

- **Bit 5 - Initiator Function**

The drive returns 0b in this field.

- **Bit 4 - Target Function**

The drive returns 1b in this field.

- **Bit 3 - Command/Data Mixed Allowed**

The drive returns 0b in this field.

- **Bit 2 - Data/Response Mixed Allowed**

The drive returns 0b in this field.

- **Bit 1 - Read XFER_RDY Disabled**

The drive returns 1b in this field.

- **Bit 0 - Write XFER_RDY Disabled**

The drive returns 0b in this field.

Table 52: PRLI/PRLO ACC response codes

Code	Description
00h	Reserved.
01h	Request executed.
02h	The target image has no resources available for establishing image pairs between the specified source and destination N_Ports. The PRLI request may be retried.
03h	Initialization is not complete for the target image. The PRLI request may be retried.
04h	The target image corresponding to the responder PA specified in the PRLI request and PRLI accept does not exist. The PRLI request shall not be retried.
05h	The target image has a predefined configuration that precludes establishing this image pair. The PRLI request shall not be retried.
06h	Request executed conditionally. Some service parameters were not able to be set to their requested state. See the service parameters response field for further details.
07h	The destination N_Port is unable to process multiple page PRLI requests. The PRLI request may be retried as a single page request.
08h-FFh	Reserved.

15.3.10 Process Logout (PRLO)

The PRLO request indicates to the responder that those process image pairs specified in the service parameter pages are being discontinued by the originator. All tasks, reservations, mode page parameters and status for the specified image pairs are set to the state they would have after a SCSI hard reset or power-on reset.

Table 53: PRLO payload

Byte	Item	Size (Bytes)
0	21h	1
1	Page Length=10h	1
2-3	Payload Length	2
4-max	Logout service parameter pages	n*16

Table 54: PRLO ACC payload

Byte	Item	Size (Bytes)
0	02h	1
1	Page Length=10h	1
2-3	Payload Length	2
4-max	Logout service parameter response pages	n*16

15.3.10.1 Process Logout Service Parameter page

Table 55: Logout Service Parameter page

Byte	Item	Size (Bytes)
0	TYPE Code = 08h	1
1	TYPE Code Extension = 00h	1
2-3	Flags	2
4-7	Originator Process Associator	4
8-11	Responder Process Associator	4
12-15	Reserved	4

The **TYPE Code** field for SCSI-FCP is 08h, and the **TYPE Code Extension** field is 00h.

The **Flags** field is as follows:

- **Bit 15 - Originator Process Associator Valid**

When set to one, the Originator Process Associator field of this Service Parameter page is valid.

- **Bit 14 - Responder Process Associator Valid**

When set to one, the Responder Process Associator field of this Service Parameter page is valid.

- **Bits 13-0 - Reserved**

The **Originator Process Associator** identifies a group of related processes (an 'image') within the originator.

The **Responder Process Associator** identifies a group of related processes (an 'image') within the responder.

15.3.10.2 Process Logout Response Service Parameter page

Table 56: Logout Response Service Parameter page

Byte	Item	Size (Bytes)	Drive Response
0	TYPE Code	1	08h
1	TYPE Code Extension	1	00h
2-3	Flags	2	0X 00h
4-7	Originator Process Associator	4	00 00 00 00h
8-11	Responder Process Associator	4	00 00 00 00h
12-15	Reserved	4	00 00 00 00h

The **TYPE Code** field for SCSI-FCP is 08h, and the **TYPE Code Extension** field is 00h.

The **Flags** field is as follows:

- **Bit 15 - Originator Process Associator Valid**

The drive returns 0b in this field.

- **Bit 14 - Responder Process Associator Valid**

The drive returns 0b in this field.

- **Bits 13-12 - Reserved**

- **Bits 11-8 - Accept Response Code**

See Table 52.

- **Bits 7-0 - Reserved**

The **Originator Process Associator** and **Responder Process Associator** fields are not used.

15.3.11 Read Link Error Status Block (RLS)

RLS requests the recipient to return the Link Error Status Block associated with the Port Identifier specified in the payload. The drive implements a Link Error Status Block for each port. When a counter overflows, it wraps back to zero. The only way to reset the Link Error Status Block is to power off the drive.

Bytes 1, 2, and 3 of the payload must be set to zero. Otherwise, the drive will respond with an LS_RJT with a reason code of “Command not supported” and a reason code explanation of “No additional explanation”.

Valid Port Identifiers are:

- 0** Return the Link Error Status Block for the same port on which the request was received.
- 1** Return the Link Error Status Block for Port A.
- 2** Return the Link Error Status Block for Port B.

Table 57: RLS payload

Byte	Item	Size (Bytes)
0-3	0F00 0000h	4
4	Reserved	1
5-7	Port Identifier	3

Table 58: RLS ACC payload

Byte	Item	Size (Bytes)
0-3	0200 0000h	4
4	Link Error Status Block	24
-		
27		

Table 59: Link Error Status block

Byte	Item	Size (Bytes)
0-3	Link Failure Count	4
4-7	Loss of sync count	4
8-11	Loss of signal count	4
12-15	Primitive Sequence Protocol error	4
16-19	Invalid Transmission Word	4
20-23	Invalid CRC Count	4

The Drive's Link Error Status block is defined as follows:

Link Failure Count: This is a count of the number of times that the port's receiver loses synchronization for a continuous period of time determined by R_T_TOV. When a link failure condition is detected by the drive, it will transmit LIP(F8) on that port.

Loss of sync count: This is a count of the number of times that the port's receiver loses synchronization.

Loss of signal count: Not implemented.

Primitive Sequence Protocol error: Not implemented.

Invalid Transmission Word: This is a count of the number of times that the port detects an invalid transmission word on its receiver.

Invalid CRC Count: This is a count of the number of frames received with invalid CRC. Only one invalid CRC is counted for each command nexus.

15.3.12 Report Node Capabilities (RNC)

The Report Node Capabilities (RNC) ELS is used to exchange node capabilities, vendor identification, and other vendor unique information. It is used to discover which document identifiers (along with their associated FC-4 protocols and profiles)

a node supports. RNC can also be used to specify which document(s) define the operating parameters between two nodes, as well as specify any additional parameters not specified during N_Port Login.

Table 60: RNC/ACC payload

Byte	Item	Size (Bytes)
0	53h for RNC, 02 for ACC	4
1	Reserved	1
2-3	Payload Length	2
4	RNC Flags	1
5-6	Reserved	2
7	VU Information Length (VU_Len)	1
8 -15	Vendor Identifier	8
16->	Capability Entry(s)	m
16+m->	Vendor Unique Information	0-128

Payload Length is a two byte unsigned integer that specifies the length of the RNC payload. The minimum length of the RNC payload is 16 bytes, and its maximum length is limited to 256 bytes. While the maximum length of the ACC payload is not defined, the drive limits its RNC_ACC to 256 bytes.

RNC Flags is a one byte field that defines options that are applicable to all Capability Entries contained in the RNC payload. The RNC Flags field is as follows:

- **Bit 7 - Select**

When this flag is set to one, the RNC_ACC payload shall contain only one Capability Entry. This Capability Entry is selected from the list of Capability Entries specified in the RNC payload. When this flag is set to zero, the RNC_ACC payload shall contain all of the Capability Entries that a node wishes to report.

The drive does not support the Select Flag and this field must be set to zero. Otherwise, the drive will respond with an LS_RJT containing a reason code of 03h “Logical error” and a reason code explanation of 2Ah “No additional explanation”.

- **Bits 6-0 - Reserved**

The **VU Information Length (VU_Len)** field is a one byte unsigned integer, which specifies the length of the Vendor Unique Information field. A maximum length of 128 bytes is supported.

The **Vendor Identifier** contains eight bytes of ASCII data, which identifies the vendor of the product.

Vendor Unique Information is defined by vendor or profile specific documentation.

15.3.12.1 Capability Entry(s)

Capability Entry(s) are used to specify standards and profiles, with which a node is compliant or supports. Nodes may also use capability entries to exchange vendor unique parameters or information.

Table 61: RNC Capability Entry

Byte	Item	Size (Bytes)
0	Flags	1
1	Document Identifier	1
2	Low Revision	1
3	High Revision	1
4-5	Reserved	0 or 2
6-7	Extension Length	0 or 2
8->	Extension	0 or n

The **Flags** field is as follows:

- **Bit 7 – Invalidate Previous**

This flag is only meaningful when the RNC Flags Select bit is set to one. Furthermore, it may only be set in the first Capability Entry in the RNC Payload. When the RNC Flags Select field and the Invalidate Previous field are both set to one, then the node issuing the RNC is requesting that the responding node invalidate this Capability Entry. All bytes of the Capability Entry marked with the Invalidate flag must match the values set in a prior RNC exchange. When the Invalidate Previous field is reset(0), the responding node selects one Capability Entry to return in the RNC_ACC.

Since the drive does not support the Select Flag (it must be set to 0), the Invalidate Previous Flag has no meaning and is ignored.

- **Bit 6 – Extended**

When this flag is set to zero, the Capability Entry shall be exactly 4 bytes long. If this flag is set to one, the length of the Capability Entry is 4 bytes plus the length of the Extension.

- **Bits 5-4 – Vendor Unique**

- **0 0** Specifies that the Document Identifier field should be interpreted according to Table 62.
- **0 1** Specifies a Vendor Unique Document Identifier.
- **1 0** Specifies a Vendor Unique Document Identifier as defined by the vendor of the N_Port receiving the RNC payload.
- **1 1** Reserved

- **Bits 3-2 – Reserved**

- **Bits 1-0 - Preference**

Preference is a two bit value that indicates the level of support or performance relative to the other capabilities supported by the node. It is used to aid a node in selecting a specific capability when multiple capabilities are supported. The Preference field has a range from 0 to 3.

- **0 0 Best**
- **0 1**

• 1 0

• 1 1 Worst

The Preference field is ignored.

The **Document Identifier** field specifies which Profile or Standard is associated with each Capability Entry. Valid Document Identifiers are listed in Table 62. If the Vendor Unique Flag is set to a value other than 00h, then the Document Identifier specifies a vendor unique capability.

Table 62: Document Identifiers

Profile or standard name	Identifier	Supported
Reserved	00h	N
FC-LE	01h	N
FC-SB	02h	N
IPI-3	03h	N
SCSI-FCP	04h	Y
FC-FP	05h	N
Reserved	06h-0Fh	N
FC-GS	10h	N
FC-FG	11h	N
FC-SW	12h	N
FC-AL	13h	Y
Reserved	14h-1Fh	N
IBM/HP/Ancor FC-PH 4.2 Deviations	20h	N
FCSI Mixed Mode SCSI Profile	21h	N
FCSI Class 2 SCSI Profile	22h	N
FCSI IP Profile	23h	N
FCSI IP Class 2 Profile	24h	N
FC-PLDA – Private Loop Direct Attach	25h	Y
FLA Fabric Loop Attach Profile	26h	Y
FCA IP Profile	27h	N
Reserved	28h-FFh	N

The **Low Revision** field defines the lowest revision of the specified document supported. The field represents a decimal revision number between 0.0 and 25.5.

The **High Revision** field defines the highest revision of the specified document supported. The field represents a decimal revision number between 0.0 and 25.5.

The **Extension Length** field is a two byte unsigned integer that specifies the number of additional bytes present in the Capability Entry. This number includes itself, the preceding reserved field, and the length of Extension Field.

The **Extension** field is used to specify any additional bit flag, parameters, or other information defined by the document associated with the Capability Entry. The drive does not currently make use of Extension field, therefore all Capability Entries returned by the drive are 4 bytes in length.

15.3.13 Re-instate Recovery Qualifiers (RRQ)

The Re-instate Recovery Qualifier (RRQ) can be sent by the Initiator to the drive, as an indication that the Recovery Qualifier (S_ID, D_ID, OX_ID, RX_ID, SEQ_ID, and SEQ_CNT) for an aborted exchange may be reused. RRQ is not required by the drive following an aborted exchange. Following an ABTS the drive allows immediate reuse of the Recovery Qualifier. The drive responds to RRQ with a valid ACC.

The following tables show the RRQ payload and the RRQ ACC payload.

Table 63: RRQ payload

Byte	Item	Size (Bytes)
0-3	Command Code = 12000000h	4
4	reserved	1
5-7	Originator S_ID	3
8-9	OX_ID	2
10-11	RX_ID	2
12 - 43	Association Header (Opt)	32

Table 64: RRQ ACC payload

Byte	Item	Size (Bytes)
0-3	ACC Code = 02000000h	4

The drive performs no protocol checking on the RRQ payload. When it recognizes an ELS with 12h in byte 0, it simply responds with an RRQ ACC.

15.3.14 Third Party Process Logout (TPRLO)

The TPRLO request is used to invalidate the operating environments between the specified image(s) at the recipient N_Port (i.e., the drive). These image pairs being invalidated are assumed to have been previously established with a Process Login (PRLI). The originator establishing the process image (with the PRLI ELS) can itself terminate the image with a Process Logout (i.e. PRLO) ELS. An originator can terminate the process pair established between another third party originator and the recipient (i.e. the drive) using TPRLO. TPRLO has the same effect as if the third party originator performed PRLO. Upon execution of a valid TPRLO ELS, the drive will set all tasks, reservations, mode page parameters, and status for the specified image pairs to the state they would have after a SCSI device reset or power-on reset.

Table 65: TPRLO payload

Byte	Item	Size (Bytes)
0	Command Code = 24h	1
1	Page Length = 10h	1
2-3	Payload Length = 14h	2
4-max	Logout service parameter pages	n*16

Table 66: TPRLO ACC payload

Byte	Item	Size (Bytes)
0	ACC Cmd Code = 02h	1
1	Page Length = 10h	1
2-3	Payload Length = 14h	2
4-max	Logout service parameter response pages	n*16

15.3.14.1 Third Party Process Logout Service Parameter page

Table 67: Logout Service Parameter page

Byte	Item	Size (Bytes)
0	TYPE Code = 08h	1
1	TYPE Code Extension = 00h	1
2-3	Flags	2
4-7	Third Party Originator Process Associator	4
8-11	Responder Process Associator	4
12	Reserved	1
13-15	Third Party Originator N_Port ID	3

The **TYPE Code** field for SCSI-FCP is 08h, and the **TYPE Code Extension** field is 00h.

The **Flags** field is as follows:

- **Bit 15 - Third Party Originator Process Associator Valid**

When set to one, the Third Party Originator Process Associator field of this Service Parameter page is valid.

- **Bit 14 - Responder Process Associator Valid**

When set to one, the Responder Process Associator field of this Service Parameter page is valid.

- **Bit 13 - Third Party Originator N_Port ID Valid**

When set to one, the Third Party N_Port ID field of this Service Parameter page is valid.

- **Bit 12 - Global Process Logout**

When set to one, all established image pairs for all N_Ports with which Process Login has been performed will be removed. This will be as if all N_Ports that had previously executed a Process Login (i.e. PRLI) with the drive had now executed the PRLO ELS. When this bit is set, only one logout parameter page shall be transmitted, and only the TYPE Code and the TYPE Code Extension fields shall have meaning.

- **Bits 11-0 - Reserved**

The **Third Party Originator Process Associator** identifies a group of related processes (an 'image') within the originator.

The **Responder Process Associator** identifies a group of related processes (an 'image') within the responder.

The **Third Party N_Port ID** specifies the N_Port associated with the image to be removed.

15.3.14.2 Third Party Process Logout Response Service Parameter page

Table 68: Logout Response Service Parameter page

Byte	Item	Size (Bytes)	Drive Response
0	TYPE Code	1	08h
1	TYPE Code Extension	1	00h
2-3	Flags	2	XX 00h
4-7	Third Party Originator Process Associator	4	00 00 00 00h
8-11	Responder Process Associator	4	00 00 00 00h
12	Reserved	1	00h
13-15	Third Party Originator N_Port ID	3	XX XX XXh

The **TYPE Code** field for SCSI-FCP is 08h, and the **TYPE Code Extension** field is 00h.

The **Flags** field is as follows:

- **Bit 15 - Third Party Originator Process Associator Valid**

The drive returns 0b in this field.

- **Bit 14 - Responder Process Associator Valid**

The drive returns 0b in this field.

- **Bit 13 - Third Party Originator N_Port ID Valid**

The drive will echo this bit from the TPRLO payload. It is not valid when the Global Process Logout flag is set to one.

- **Bit 12 - Global Process Logout**

The drive will echo this bit from the TPRLO payload.

- **Bits 11-8 - Accept Response Code**

See Table 52.

- **Bits 7-0 - Reserved**

The **Originator Process Associator** and **Responder Process Associator** fields are not used.

The drive will echo the **Third Party N_Port ID** field from the TPRLO request payload when the **Third Party Originator N_Port ID Valid Flag** is set to one and the **Global Process Logout Flag** is set to zero.

15.3.15 Request Node Identification Data (RNID)

The Request Node Identification Data ELS provides a mechanism for a node to acquire Node Identification Data from other nodes in a Fibre Channel fabric. This function is normally used by nodes that wish to determine the topology of the network to which they are attached.

The normal response to a RNID Req is RNID ACC from the drive containing the requested Node Identification information. RNID is available prior to LOGIN.

Table 69: “RNID payload” shows the RNID payload.

Table 69: RNID payload

Item	Size (Bytes)
hex '78 00 00 00'	4
Node Identification Data Format	1
Reserved	3

The **Node Identification Data Format** specifies the format of Node Identification Data returned from the drive. Acceptable values for this field are 00h (Common Node Identification Data only) or DFh (Topology Discovery data). If this field is set to any other value, the drive will return a LS_RJT containing a reason code of 03h “Logical Error” and a reason code explanation of 2Ah “No additional explanation”.

Table 70: “RNID accept payload” shows the RNID ACC payload.

Table 70: RNID accept payload

Item	Size Bytes
hex '02 00 00 00'	4
Node Identification Data Format	1
Common Node-Identification-Data Length	1
Reserved	1
Specific Node-Identification-Data Length	1
Common Node Identification Data	0 or 16
Specific Node Identification Data	0-max

The **Node Identification Data Format** specifies the format of Node Identification Data returned from the drive. The drive sets this field to the same value that was in the RNID Request.

The **Common Node-Identification-Data Length** specifies the length of the Common Node-Identification-Data. This field is set to 10h.

The **Specific Node-Identification-Data Length** specifies the length of the Specific Node-Identification-Data. This field is set to 34h if the Node Identification Data Format was set to DFh (Topology Discovery Data) or 0h otherwise.

Common Node-Identification-Data specifies the port's 8-byte Node_Name and 8-byte Port_Name. The format is defined in Table 71: "Common Node Identification Data".

Table 71: Common Node Identification Data

Format	Size (Bytes)
N_Port_Name	8
Node_Name	8

The **Specific Node-Identification-Data** specifies Topology Discovery Data whose format is described in Table 72: "Topology Discovery Specific Node Identification Data". This Data is only returned when the **Node Identification Data Format** is set to DFh "Topology Discovery Data".

Table 72: Topology Discovery Specific Node Identification Data

Format	Size (Bytes)
Global ID	16
Unit Type	4
Physical Port Number	4
Number of Attached Nodes	4
IP Version	2
UDP Port Number	2
IP Address	16
Reserved	2
Topology Discovery Flags	2

The **Global ID** field is a World Wide unique name whose format is identical to the Common Node-Identification-Data field as specified in the table.

The **Unit Type** field specifies the type of device returning the Node-Identification-Data. Valid values for this field are shown in Table 73: "Topology Discovery Unit Type".

Table 73: Topology Discovery Unit Type

Value - hex	Type
'00 00 00 00'	Reserved
'00 00 00 01'	Unknown
'00 00 00 02'	Other (none of the following)
'00 00 00 03'	Hub
'00 00 00 04'	Switch
'00 00 00 05'	Gateway
'00 00 00 06'	Converter
'00 00 00 07'	HBA
'00 00 00 08'	Proxy-agent
'00 00 00 09'	Storage device (drive, CD, tape, etc.)
'00 00 00 0A'	Host
'00 00 00 0B'	Storage subsystem (raid, library, etc.)
'00 00 00 0C'	Module (subcomponent of a system)
'00 00 00 0D'	Software driver
'00 00 00 0E' - 'FF FF FF FF'	Reserved

The drive returns a value of 9h in this field.

The **Physical Port Number** specifies the port that received the RNID REQ. Valid values are 00h if the RNID was received on port A and 01h if the RNID was received on port B.

The **Number of Attached Nodes** field specifies the number of nodes attached to the node returning the RNID ACC. Since the drive does not perform any topology discovery, it sets this field to 0h.

The **IP Version** field specifies the level of IP supported. Since the drive does not support the IP protocol, it sets this field to 0h.

The **UDP Port** field specifies the numerical value that identifies a port using the User Datagram Protocol. Since the drive does not support UDP, it sets this field to 0h.

The **IP Address** specifies the IP address of the node. Since the drive does not support IP, this field is set to 0h.

The **Topology Discovery Flags (TDF)** are defined in Table 74: “Topology Discovery Flags” .

Table 74: Topology Discovery Flags

Byte	BIT								Default request values
	7	6	5	4	3	2	1	0	
0	Reserved						TDF		000h

The **Topology Discovery Support** flag signals that the node supports further Topology Discovery inquiries. This flag is set to 0 by the drive.

The **Loop Position Valid** flag indicates that multiple Node Identification Data records are reported in the order detected by a Loop Position Report Primitive. The drive returns only a single Node Identification Data record and sets this flag to 0.

15.3.16 Report Port Speed Capabilities (RPSC)

The Report Port Speed Capabilities ELS provides a mechanism for a port to report its current and potential link operating speeds. The response to the RPSC payload is an LS_ACC response which indicates speed capabilities and current operating speeds of each port.

Table 75: RPSC payload

Byte	Item	Size (Bytes)
0-3	Command Code = 7D000000h	4

Table 76: RPSC ACC payload

Byte	Item	Size (Bytes)
0	ACC Cmd Code = 02h	1
1	Reserved = 0	1
2-3	Number of Entries = 0002h	2
4-5	Port 1 Speed Capabilities = E000h	2
6-7	Port 1 Operating Speed	2
8-9	Port 2 Speed Capabilities = E000h	2
10-11	Port 2 Operating Speed	2

The **Speed Capabilities** fields indicate that the drive is capable of 1Gb/sec., 2 Gb/sec., and 4 Gb/sec. operations.

The **Operating Speed** fields indicate the port's current operating speed as follows: 8000h for 1Gb/sec., 4000h for 2 Gb/sec., or 2000h for 4 Gb/sec.

15.4 Common Fibre Channel Services

Common Fibre Channel Service Request uses the Common Transport (CT) Interface to communicate with FC Fabric Servers. The CT interface provides several Information Units, (CT_IU)s, which provide the transport mechanism between service requestors and their servers.

In CT_IU frames, the TYPE field in the frame header is set to 20h (Fibre Channel Service). R_CTL is either 02h (FC-4 Device_Data, Unsolicited Control) for a Request or 03h (FC-4 Device_Data, Solicited Control) for a Reply.

Each CT_IU payload starts with a 16 byte CT Header as shown in Table 77.

Note: Common FC Services are used to communicate with the Fabric Name Server and are only issued by the drive when it is operating as a Public Loop Device.

Table 77: Payload of a CT Header

Byte	Item	Size (Bytes)
0	FC_CT	1
1-3	IN_ID	3
4	FCS_Type	1
5	FCS_Subtype	1
6	Options	1
7	Reserved	1
8-9	Command/Response Code	2
10-11	Maximum/Residual Size	2
12	Reserved	1
13	Reason Code	1
14	Reason Code Explanation	1
15	Vendor Unique	1

The **FC_CT** field contains the FC_CT revision. The drive always sets this value to 0x01.

IN_ID is a reserved field that may be used to carry the S_ID of the original requestor between different servers in the fabric. Its value in the FC_ACC/FC_FJT frame may be non zero if it was used.

The **FCS_Type** indicates which FC Service is requested and the **FCS_Subtype** specifies which server will process the request. The drive does not check the FCS_Type/FCS_Subtype values returned by the Name Server in the FS_ACC/FS_RJT.

The **Options** field specifies various options used during FCS processing. The drive sends a value of 0x00, which specifies that this request must complete before another can be attempted.

The **Command/Response Code** field identifies specific Fibre Channel Service Information Units (FS_IUs). The Command/Response Code also determines the type of FS_IU. Command codes are as shown in Table 78.

Table 78: Command/Response Codes

Value	Description
0000h	Non-(FS_IU)
0001h - 7FFFh	Fibre Channel Service Requests (FS_REQ)
8001h	Fibre Channel Service Reject (FS_RJT)
8002h	Fibre Channel Service Accept (FS_ACC)
other values	Reserved

The drive supports only the Command code of 0217h, which specifies a RFT_ID request.

The **Maximum/Residual Size** field is used by the sender of an FS_REQ to indicate the maximum number of bytes it can accept in an FS_ACC payload. If the size of the responder's FS_ACC payload is greater than this value, the responder will

transfer only the number of bytes requested and set the Maximum/Residual field in the FS_ACC to the number of residual bytes that were not transferred.

15.4.1 Register FC-4 Types (RFT_ID)

The drive uses RFT_ID to register its FC-4 type, (SCSI-FCP), with the Fabric Name Server. An RFT_ID request is sent to the well-know address FFFFFCh (Fabric Name Server) and must be attempted after a successful FLOGI.

Table 79: RFT_ID payload

Byte	Item	Size (Bytes)
0 - 15	CT_IU Header	16
16	Reserved	1
17 - 19	S_ID of sending NL_Port	3
20 - 51	Supported FC-4 types bit map: byte 22 = 0x01, all others = 0x00	32

As with all Common Transport requests, the payload of an RFT_ID starts with the CT_IU header. The FCS_Type value of the CT_IU Header is set to FCh to indicate a Directory Service Request, while the FCS_Subtype value is set 02h to indicate that the request is directed to the Name Server. The Command code is set to 0217h to specify an RFT_ID request. For an RFT_ID_ACC/RJT, the Command/Response code is set to 8001h for a RFT_ID_RJT and 8002h for RFT_ID_ACC. If the response was a FA_RJT, the Reason Code, Reason Code Explanation, and Vendor Unique fields would be set to indicate which error condition occurred. The drive does not check any other fields in an FS_ACC/RJT payload.

The **Supported FC-4 Bitmap** indicates which FC-4 protocols the RFT_ID requestor supports. The drive sets the eighth bit of this bitmap to indicate that it supports FC-4 type 08h (SCSI-FCP).

Table 80: RFT_ID ACC/RJT payload

Byte	Item	Size (Bytes)
0 - 15	CT_IU Header	16

Table 81: FS_RJT Reason Codes

Encoded value	Description
0000 0001	Invalid command code
0000 0010	Invalid version level
0000 0011	Logical error
0000 0100	Invalid IU size
0000 0101	Logical busy
0000 0111	Protocol error
0000 1001	Unable to perform command request
0000 1011	Command not supported
others	Reserved
1111 1111	Vendor Unique Error

Table 82: FS_RJT Reason Explanations

Encoded value	Description
00	No additional explanation
01	Port Identifier not registered
02	Port Name not registered
03	Node Name not registered
04	Class of Service not registered
05	IP address not registered
06	Initial Process Associator not registered
07	FC-4 TYPEs not registered
08	Symbolic Port Name not registered
09	Symbolic Node Name not registered
0A	Port Type not registered
10	Access Denied
11	Unacceptable Port Identifier
12	Data base empty
Others	Reserved

15.5 FC-AL timers

Table 83: FCAL timer values

Timer	Description	Value (PLDA/FLA)*
AL_TIME	Arbitrated Loop Timeout Value	15 ms
LIS_HOLD_TIME	Loop Initialization Sequence Hold Time	1 ms
R_T_TOV	Receiver Transmitter Timeout Value	100 ms
E_D_TOV	Error Detect Timeout Value	2 sec./2 sec.*
R_A_TOV	Resource Allocation Timeout Value	2 sec./10 sec.*
RR_TOV	Resource Recovery Timeout Value	2 sec.
LP_TOV	Loop Timeout Value	2 sec.

The drive uses FCAL timers as specified in PLDA 2.1 as required in the standard.

- **AL_TIME**

AL_TIME represents two times the worst case round trip latency for a very large loop.

- **LIS_HOLD_TIME**

LIS_HOLD_TIME is the maximum amount of time between when a node receives a Loop Initialization Sequence until it forwards it to the next node.

- **R_T_TOV**

R_T_TOV is used by the receiver logic to detect a Link Failure. A Link Failure is defined as loss of synchronization for a period greater than R_T_TOV.

- **E_D_TOV**

E_D_TOV is the minimum time a port shall wait for the next expected frame in a sequence before detecting an error.

- **R_A_TOV**

R_A_TOV is defined by the PLDA to be two distinct timers.

- **R_A_TOV_{SEQ_QUAL}** defines the minimum time a port shall wait before reuse of the sequence qualifiers SEQ_CNT and SEQ_ID. The PLDA defines the value of R_A_TOV_{SEQ_QUAL} to be zero seconds.

Twice the value of R_A_TOV_{ELS} defines the minimum time a port shall wait for the response to an Extended Link Service Request. The PLDA defines this value to be two seconds. Therefore a port must wait twice this value (four seconds) before timing out an ELS.

Note: In the FLA, there is no split definition of the R_A_TOV value. One timer value is used for both R_A_TOV_{SEQ_QUAL} and R_A_TOV_{ELS}.

- **RR_TOV**

RR_TOV is the minimum time a target shall wait for an Initiator to perform exchange authentication following LIP. If this timer expires, the Target will implicitly log out the Initiator and free up the resources associated with that timer. An RR_TOV timer is maintained for each Initiator that has logged in.

- **LP_TOV**

LP_TOV is used to keep a Loop from deteriorating due to protocol errors or lost Ordered Sets.

LP_TOV is also used during initialization and to reset the fairness window.

15.5.1 Link Failure

Link Failure is defined when a receiver has continuously detected loss of synchronization for a period of R_T_TOV. When this occurs, the drive will transmit LIP(F8) on that port.

15.6 Invalid frame delimiter

If an invalid frame delimiter is received for a:

- FCP_DATA frame, the exchange shall terminate with a CHECK CONDITION status. The resulting Sense data will have the Sense Key set to *Aborted command* and the Additional Sense Code set to *Data Phase Error*.
- non-FCP_DATA frame, the frame shall be discarded and ignored.

16.0 SCSI-FCP

This section describes the drive's implementation of SCSI-FCP. SCSI-FCP is the FC-4 mapping recommended by the drive Profile. It maps the ANSI SCSI protocol onto the FC-PH functions.

Note: PLDA 2.1 specifies: "Reserved FC-PH fields are not required to be checked for zeroes. Validity bits set to 0 remove any requirement to check the corresponding field for zeroes (e.g., if F_CTL bit 3=0, receiving N_Ports are not required to verify that the parameter field in word 5 of the frame header contains zeroes)." As such, the drive does not validate 1) reserved FC fields or 2) fields that are not reserved but are not valid for the current frame (as the example above with F_CTL bit 3). This does not apply to any reserved field checking and testing within the FCP_CDB. These fields are checked as per ANSI SCSI requirements.

16.1 Terminology

The SCSI Architecture Model (SAM) defines a new SCSI vocabulary in order to remain independent from physical protocol and interconnect. Common SCSI-2 terms have been replaced and new ones introduced.

As might be expected, the terminology used to describe SCSI-FCP is a problem. It needs elements of SAM and Fibre Channel. Most of the Fibre Channel terms have been introduced in 1.0, "FC-AL Attachment", and some essential SAM terms are covered here. However, SAM is still recommended reading!

SCSI-FCP describes all communication in terms of Fibre Channel Information Units (IUs). FC-PH defines these simply as "sequences that have special meaning to the FC-4" (i.e. SCSI-FCP).

IUs are used to send commands, data, and status; bus phases are part of parallel SCSI and no longer exist. Execution of a command requires several IUs to pass between the target and Initiator so an I/O Process (IOP) equates to a Fibre Channel exchange. In fact, the exchange ID is used as the command tag.

An IOP is represented in the Target by a Task. The Initiator uses Task Management functions to control execution of the task in the Target. IUs are used to transfer Task Management functions.

SAM defines a Target as consisting of a Task Manager and one or more Logical Units. The Task Manager handles all the Task Management functions and the logical unit handles commands.

The logical unit consists of a Target and a Task Set. The Target actually executes the commands and the task set is simply what used to be known as the Command Queue.

SCSI-FCP recognizes that Targets and Initiators may simply be software procedures and therefore calls them Processes.

16.2 Information Units

SCSI-FCP defines a number of IUs that are used to describe the mapping of SAM Device and Task Management functions. The drive Profile defines an 'FCP Feature Set', which is a subset of those IUs.

This section defines the IUs implemented by the drive (see Table 84).

Table 84: Information Units (IUs)

IU	SAM primitive	Data block		F/M/L	SI	M/O
		CAT	Content			
T1	Command Request	6	FCP_CMND	F	T	M
T6	Data Out action	1	FCP_DATA	M	T	M
I1	Data delivery request	5	FCP_XFER_RDY (WRITE)	M	T	M
I3	Data In action	1	FCP_DATA	M	H	M
I4	Response	7	FCP_RSP	L	T	M

Note:

Key:

- SAM = SCSI-3 Access Method
- IU = Information Unit
- CAT = Information category of Device_Data frame
- F/M/L = First/Middle/Last IU of Sequence
- SI = Sequence Initiative
- H = Hold Sequence Initiative
- T = Transfer Sequence Initiative
- M/O = Mandatory/Optional

16.2.1 FCP_CMND

The FCP_CMND IU carries either a SCSI command to be executed or a Task Management function to be performed.

Table 85: FCP_CMND payload

Byte	Field	Description	Size (Bytes)
0 -7	FCP_LUN	Logical Unit Number	8
8-11	FCP_CNTL	Control Field	4
12 - 27	FCP_CDB	SCSI Command Descriptor Block	16
28-31	FCP_DL	Data Length	4

16.2.1.1 FCP_LUN

The FCP_LUN field identifies the logical unit number within the Target. The drive is a single LUN with address 0000 0000 0000 0000h.

16.2.1.2 FCP_CNTL

The FCP_CNTL field contains a number of control flags.

Table 86: FCP_CNTL field

Byte	Description	Size (Bytes)
0	Reserved	1
1	Task Codes	1
2	Task Management function flags	1
3	Execution management codes	1

The **Task Codes** field contains the Task Attributes, as shown in Table 87. The Task Attributes are described in 16.3, “Task Attributes”, on page 99.

Table 87: Task Attribute values

Value	Attribute
xxxx x000b	Simple_Q
xxxx x001b	Head_of_Q
xxxx x010b	Ordered_Q
xxxx x100b	ACA_Q (not supported)
xxxx x101b	Untagged

The **Task Management Function flags** are used to request Task Management functions, as shown in Table 88. The Task Management functions are described in 16.4, “Task Management functions”, on page 100.

Table 88: TMF flag values

Value	Function
1000 0000b	Terminate Task
0100 0000b	Clear ACA (not supported)
0010 0000b	Target Reset
0000 0100b	Clear Task Set
0000 0010b	Abort Task Set

The **Execution management codes** field contains flags, as follows:

Bits 7-2 - Reserved

Bit 1 - Read Data

When set to one, this flag indicates that the Initiator expects FCP_DATA IUs for the task to be in the direction opposite to the direction of the FCP_CMND IU. This is a SCSI READ type operation.

Bit 0 - Write Data

When set to one, this flag indicates that the Initiator expects FCP_DATA IUs for the task to be in the same direction as the FCP_CMND IU. This is a SCSI WRITE type operation.

16.2.1.3 FCP_CDB

The **FCP_CDB** field contains the SCSI CDB to be executed by the addressed logical unit. This field is ignored if any of the Task Management function flags are set.

This is a 16-byte field. Bytes beyond the end of the CDB are ignored by the Target and may have any value.

16.2.1.4 FCP_DL

The **FCP_DL** field contains a count of the greatest number of data bytes expected to be transferred by execution of the SCSI CDB. An FCP_DL field of zero indicates that no data transfer is expected and that no FCP_XFER_RDY or FCP_DATA IUs shall be transferred.

16.2.2 FCP_XFER_RDY

The FCP_XFER_RDY IU indicates that the Target is prepared to perform all or part of the data transfer for a command. During WRITE operations, the FCP_XFER_RDY IU indicates the amount of data that the Target expects from the Initiator. Since the Target has planned buffer resources based on that amount of data, the Initiator is expected to provide exactly the amount requested.

Table 89: FCP_XFER_RDY payload

Byte	Field	Description	Size (Bytes)
0-3	DATA_RO	Relative Offset	4
4-7	BURST_LEN	Length of FCP_DATA IU that follows	4
8-11	Reserved		4

The **DATA_RO** field indicates the contents of the RLTV_OFF field for the first data byte of the next FCP_DATA IU. The RLTV_OFF field is part of the frame header.

The **BURST_LEN** field indicates the amount of buffer space prepared for the next FCP_DATA IU and requests an IU of that exact length.

16.2.3 FCP_DATA

SCSI data transfers may be performed by one or more data delivery requests, each one performing a transfer no longer than the maximum burst length defined by the parameters of the disconnect/reconnect mode page.

16.2.4 FCP_RSP

The content of the FCP_RSP IU is as shown in Table 90.

Table 90: FCP_RSP payload

Byte	Field	Description	Size (Bytes)
0-7	Reserved		8
8-11	FCP_STATUS	Field Validity and SCSI Status	4
12-15	FCP_RESID	Residual Count	4
16-19	FCP_SNS_LEN	Length of FCP_SNS_INFO field	4
20-23	FCP_RSP_LEN	Length of FCP_RSP_INFO field	4
24->	FCP_RSP_INFO	FCP Response Information	m
24+m->	FCP_SNS_INFO	SCSI Sense Information	n

16.2.4.1 FCP_STATUS

The FCP_STATUS field is normally zero upon successful completion of an IOP.

Table 91: FCP_STATUS field

Byte	Description	Size (Bytes)
0-1	Reserved	2
2	Flags	1
3	SCSI Status byte	1

The **Flags** field contains the following:

Bits 7-4 - Reserved

Bit 3 - FCP_RESID_UNDER

When set to one, this flag indicates that the FCP_RESID field is valid and contains a count of the number of bytes that were expected but not received.

Bit 2 - FCP_RESID_OVER

When set to one, this flag indicates that the FCP_RESID field is valid and contains a count of the number of bytes that could not be transferred because FCP_DL was not sufficient.

Bit 1 - FCP_SNS_LEN_VALID

When set to one, this flag indicates that the FCP_SNS_LEN field is valid and contains a count of the number of bytes in the FCP_SNS_INFO field.

Bit 0 - FCP_RSP_LEN_VALID

When set to one, this flag indicates that the FCP_RSP_LEN field is valid and contains a count of the number of bytes in the FCP_RSP_INFO field.

The **SCSI Status byte** field is defined in 18.0, “SCSI Status Byte”, on page 271.

16.2.4.2 FCP_RESID

This field contains a count of the number of residual data bytes that were not transferred for this SCSI command.

16.2.4.3 FCP_SNS_LEN

This field contains a count of the number of valid bytes in the FCP_SNS_INFO field.

16.2.4.4 FCP_RSP_LEN

If the FCP_RSP_LEN_VALID flag in the FCP_STATUS field is set to 1b, this field contains a count of the number of valid bytes in the FCP_RSP_INFO field.

Valid values are 0, 4, and 8.

16.2.4.5 FCP_RSP_INFO

This field contains information describing only the protocol failures detected during the execution of an IOP.

Table 92: FCP_RSP_INFO field

Byte	Description	Size (Bytes)
0-2	Reserved	3
3	RSP_CODE	1

The content of the **RSP_CODE** field is defined below.

Table 93: RSP_CODE definitions

Code	Description
00h	No Failure (Function Complete)
01h	FCP_DATA length different from BURST_LEN
02h	FCP_CMND Fields Invalid
03h	FCP_DATA_RO mismatch with FCP_XFER_RDY DATA_RO
04h	Function Rejected
05h	Service Delivery or Target Failure
06-FFh	Reserved

16.2.4.6 FCP_SNS_INFO

This field contains the Sense information specified by SCSI (see 20.1, “SCSI Sense Data Format”, on page 309).

FCP_SNS_INFO is only returned if the SCSI Status byte in the FCP_STATUS field is set to CHECK CONDITION status. Refer to 16.5.3, “Autosense”, on page 103 for more details.

16.3 Task Attributes

Task Attributes are specified in the Task Codes field of the FCP_CMND IU. They apply only to the SCSI command contained within the FCP_CDB field and are ignored if any of the Task Management function flags are set.

A task shall have one of the following attributes:

Simple Queue

This attribute specifies that the task shall be accepted into the task set and executed after tasks with the Head of Queue attribute and tasks with the Ordered attribute that were received earlier. The order of execution, with respect to other tasks with the Simple attribute, is determined by the Queue Algorithm currently in effect. Refer to the Queue Algorithm Modifier field in 17.10.9, “Mode Page 0A (Control Mode Page Parameters)”, on page 169.

Head of Queue

This attribute specifies that the task shall be accepted into the task set and executed next. Successive tasks received with Head of Queue attribute will be executed in LIFO order.

Ordered Queue

This attribute specifies that the task shall be accepted into the task set and executed in the order received. All tasks received earlier shall complete before this task. All tasks received later shall complete after this task, except for tasks received with Head of Queue attribute.

Untagged

This attribute specifies that the task shall be accepted into the task set according to the rules for an untagged task. Only one untagged task can exist for each logical unit/Initiator pair. A second untagged task for the same pair is treated as an overlapped command. SCSI-FCP commands are inherently tagged with the OXID/RXID specified in the frame header. Therefore, commands sent with the untagged task attribute are handled as if they had the Ordered Queue attribute.

Note: The uncoupled nature of FCAL makes it impossible for the drive to perform command overlap checking. (See Section 21.6.3 “Overlapped Commands”)

16.4 Task Management functions

Task Management functions allow an Initiator to explicitly control the execution of one or more Tasks. SAM defines the following:

- Abort Task (Implemented as ABTS BLS)
- Abort Task Set
- Reset LUN
- Clear ACA (not supported)
- Clear Task Set
- Target Reset
- Terminate Task

Note: The following description applies to all Task Management functions except Abort Task, which is described separately in 16.4.1, “Abort Task (Implemented as ABTS BLS)”, on page 101.

Task Management functions are transmitted in a new exchange using a T1 FCP_CMND IU (see 16.2.1, “FCP_CMND”, on page 95). The Task Management function flags in the IU specify the required function. If any flag is set, all CDB related fields of the IU are ignored (FCP_CDB, FCP_DL, the Task Codes and Execution Management fields of FCP_CNTL). Only one Task Management function flag may be set.

The Target responds to a Task Management function with an I4 FCP_RSP IU (see 16.2.4, “FCP_RSP”, on page 97). The RSP_CODE in the FCP_RSP_INFO field shall be as shown in Table 94 and all other fields shall be zeroes.

Table 94: Task Management function RSP_CODE definitions

Code	Description
00h	No Failure (Function Complete)
04h	Function Rejected
05h	Service Delivery or Target Failure

16.4.1 Abort Task (Implemented as ABTS BLS)

The Abort Task function is performed using the FC-PH link management functions. Specifically, the Abort Sequence (ABTS) Basic Link Service (BLS). Refer to 15.2.1, “Abort sequence (ABTS)”, on page 47 for a detailed description of this BLS.

Abort Task causes the Target to abort the specified task, if it exists. Previously established conditions such as Mode parameters and reservations are not affected.

The Initiator performs the Abort Task by using the FC-PH Recovery Abort Protocol. This protocol recovers any resources associated with the exchange that is being terminated.

The Initiator shall also use the recovery abort protocol for each open exchange following receipt of an FCP_RSP with a RSP_CODE of No Failure (Function Complete) to Abort Task Set, Clear Task Set, or Target Reset Task Management function.

The recovery abort protocol is as follows:

- The Initiator generates an ABTS sequence. This may be done regardless of whether or not the Initiator has sequence initiative. In the ABTS frame:
 - **SEQ_ID** is set to the SEQ_ID of an open sequence at the Initiator. If no sequence is open, any SEQ_ID not currently in use between the Initiator and the Target may be used.
 - **SEQ_CNT** is set to one more than the SEQ_CNT of the last frame transmitted in the open sequence. If no sequence is open, SEQ_CNT is zero.
 - **OX_ID** is set to the OX_ID assigned by the Initiator to the task that is being aborted.
 - **RX_ID** is set to FFFFh if no FCP_XFR_RDY (WRITE) or READ data (READ) has been received from the Target. Otherwise, RX_ID is set to the RX_ID assigned by the Target to the task that is being aborted.
 - **F_CTL** is set for Sequence Context = Initiator.
- The Target may reject the ABTS with a BA_RJT frame only if
 1. the Target has assigned an RX_ID to an OX_ID in a previous frame, and
 2. the ABTS has an unknown OX_ID/RX_ID combination.

The reason code in the BA_RJT is Logical error and the reason code explanation is Invalid OX_ID-RX_ID combination. The F_CTL field is set to indicate that this is the last frame of the exchange.

- The target accepts the ABTS with a BA_ACC frame. In the frame header:
 - **OX_ID** is set from the OX_ID in the ABTS
 - **RX_ID** is set from the RX_ID in the ABTS
 - **F_CTL** is set to indicate that this is the last frame of the exchange (L_S = 1b) and Sequence Context =Recipient.

In the frame payload:

- **SEQ_ID** Validity is set to 00h
- **SEQ_ID** is ignored by the recipient and may therefore be any value
- **OX_ID** is set from the OX_ID in the ABTS
- **RX_ID** is set from the RX_ID in the ABTS
- **Low SEQ_CNT** is set to 0000h
- **High SEQ_CNT** is set to FFFFh

16.4.2 Abort Task Set

Abort Task Set causes the Target to abort all tasks in the task set that were created by the Initiator. The action is equivalent to receiving a series of Abort Task requests. Previously established conditions such as Mode parameters and reservations are not affected.

16.4.3 Terminate Task

The drive does not support Terminate Task. An FCP_RSP with a RSP_CODE of Function Rejected will be returned.

16.4.4 Clear ACA

The drive does not support Clear ACA. An FCP_RSP with a RSP_CODE of Function Rejected will be returned.

16.4.5 Target Reset

The Target Reset Task Management function causes the Target to execute a hard reset, as defined by SAM. This means:

1. Abort all tasks for all Initiators.
2. Release any device reservation.
3. Return all internal states to their initial power-on and default values, as established by PRLI.
4. Set a unit attention condition for all Initiators.

Target Reset does not affect any login state.

16.4.6 Clear Task Set

Clear Task Set causes the Target to abort all tasks in the task set. The action is equivalent to receiving a series of Abort Task requests from all Initiators.

A unit attention condition shall be generated for all other Initiators with tasks in the task set. The Additional Sense Code shall be Commands Cleared by another Initiator.

Previously established conditions such as Mode parameters and reservations are not affected.

16.4.7 Reset LUN

The Reset LUN Task Management function causes the drive to execute a hard reset with the same actions as Target Reset. See section 18.4.5.

16.5 Miscellaneous

16.5.1 Tags

Using SCSI-FCP, an I/O Process equates to a Fibre Channel Exchange. All IOPs are therefore implicitly tagged by the Fully Qualified Exchange ID (FQXID).

The FQXID is the 80-bit concatenation of the S_ID, D_ID, OX_ID and RX_ID from any frame in the exchange.

16.5.2 Auto-Contingent Allegiance (ACA)

In parallel SCSI-2, a CHECK CONDITION generates sense data for the faulted initiator (the one that got the CHECK CONDITION) and either suspends or aborts all commands queued by that Initiator. Execution of the next command from the faulted Initiator clears the sense data and allows queued commands to resume. This is 'Contingent Allegiance' (CA).

CA requires an interlocked interface to work. In a non-interlocked serial interface, the 'next' command could already be on its way before the faulted Initiator receives the CHECK CONDITION and sense data could be inadvertently lost.

In SCSI-3, ACA replaces the SCSI-2 Contingent Allegiance and Extended Contingent Allegiance conditions.

ACA is a condition that is created within the Target task set whenever CHECK CONDITION is returned. While it exists, all tasks in the task set are blocked (i.e., the queue is frozen). New tasks from Initiators other than the faulted Initiator are not entered into the task set and are completed with a status of ACA ACTIVE.

New tasks from the faulted Initiator are entered into the task set under two different conditions, depending on the state of the NACA bit in the CDB Control Byte of the faulting command.

1. If the NACA bit was set to 0b, the previously described SCSI-2 CA rules apply.
2. If the NACA bit was set to 1b, the new task must have the ACA attribute and there must be no other task in the task set with the ACA attribute. If these conditions are not met, the task is completed with a status of ACA ACTIVE.

Note: Only NACA = 0b is supported.

16.5.3 Autosense

Autosense is the automatic return of sense data upon completion of the task. It is equivalent to an explicit REQUEST SENSE command being executed immediately after *Check Condition* status is returned.

The sense data is returned in the I4 FCP_RSP IU. The FCP_SNS_LEN field indicates how many bytes of valid sense are contained in the FCP_SNS_INFO field.

17.0 SCSI Command Set

Summaries of the SCSI commands supported by the drive are listed below. O = optional, M = mandatory

Table 95: SCSI Commands Supported

Type	Code	Description
M	04h	FORMAT UNIT (04), page 108
M	12h	INQUIRY (12), page 112
O	4Ch	LOG SELECT (4C), page 124
O	4Dh	LOG SENSE (4D), page 127
O	15h	MODE SELECT (15), page 152
O	55h	MODE SELECT (55), page 153
O	1Ah	MODE SENSE (1A), page 154
O	5Ah	MODE SENSE (5A), page 178
O	5Eh	PERSISTENT RESERVE IN (5E), page 179
O	5Fh	PERSISTENT RESERVE OUT (5F), page 182
O	34h	PRE-FETCH (34), page 187
M	08h	READ (6) - (08), page 188
M	28h	READ (10) - (28), page 189
O	A8h	READ (12) - (A8), page 191
O	88h	READ (16) - (88), page 192
O	7Fh/09h	READ (32) - (7F/09), page 193
O	3Ch	READ BUFFER (3C), page 195
M	25h	READ CAPACITY (10) - (25), page 199
O	9Eh/10h	READ CAPACITY (16) (9E/10), page 201
O	37h	READ DEFECT DATA (37), page 203
O	B7h	READ DEFECT DATA (B7), page 207
O	3Eh	READ LONG (3E), page 209
O	07h	REASSIGN BLOCKS (07), page 210
O	1Ch	RECEIVE DIAGNOSTICS RESULTS (1C), page 212
M	17h	RELEASE (17), page 216
O	57h	RELEASE (57), page 217
O	A3h/05h	REPORT DEVICE IDENTIFIER (A3/05), page 218
O	A0h	REPORT LUNS (A0), page 220
O	A3h/0Ch	REPORT SUPPORTED OPERATION CODES (A3/0C), page 221
O	A3h/0Dh	REPORT SUPPORTED TASK MANAGEMENT FUNCTIONS (A3/0D), page 226
M	03h	REQUEST SENSE (03), page 227
M	16h	RESERVE (16), page 228
O	56h	RESERVE (56), page 229
O	01h	REZERO UNIT (01), page 230
M	1Dh	SEND DIAGNOSTIC (1D), page 232
O	A4h/06h	SET DEVICE IDENTIFIER (A4/06), page 237
O	1Bh	START STOP UNIT (1B), page 238
O	35h	SYNCHRONIZE CACHE (10) - (35), page 239
O	91h	SYNCHRONIZE CACHE (16) - (91), page 240

M	00h	TEST UNIT READY (00), page 241
O	2Fh	VERIFY (2F), page 244
O	AFh	VERIFY (12) - (AF), page 246
O	AFh	VERIFY (16) - (8F), page 247
O	7Fh/0Ah	VERIFY (32) - (7F/0A), page 248
M	0Ah	WRITE (6) - (0A), page 250
M	2Ah	WRITE (10) - (2A), page 251
O	AAh	WRITE (12) - (AA), page 253
O	8Ah	WRITE (16) - (8A), page 254
O	7Fh/0Bh	WRITE (32) - (7F/0B), page 255
O	2Eh	WRITE AND VERIFY (10) - (2E), page 257
O	AEh	WRITE AND VERIFY (12) - (AE), page 258
O	8Eh	WRITE AND VERIFY (16) - (8E), page 259
O	7Fh/0Ch	WRITE AND VERIFY (32) - (7F/0C), page 260
O	3Bh	WRITE BUFFER (3B), page 261
O	3Fh	WRITE LONG (3F), page 264
O	41h	WRITE SAME (41), page 266
O	93h	WRITE SAME (16) - (93), page 267
O	7Fh/0Dh	WRITE SAME (32) - (7F/0D), page 268
O	42h	UNMAP (42)

17.1 SCSI Control Byte

The Control Byte is the last byte of every CDB. The format of this byte is shown below.

Table 96: SCSI Control Byte

BIT							
7	6	5	4	3	2	1	0
VU = 0		Reserved = 0				FLAG	LINK

VU

VU stands for Vendor Unique.

FLAG**

If Link is zero, Flag must also be zero. If Link is one, Flag may also be one. Typically this bit is used to cause an interrupt in the Initiator between linked commands.

LINK**

This bit is set to one to indicate that the Initiator desires an automatic link to the next command upon successful completion of the current command.

Note: * - The drive ignores the link bit and flag bit in the CDB.

17.2 Abbreviations

These abbreviations are used throughout the following sections:

LUN Logical Unit Number. An encoded three bit identifier for the logical unit.

VU Vendor Unique bits

LBA Logical Block Address

RSVD Reserved

MSB Most Significant Byte

LSB Least Significant Byte

17.3 Byte ordering conventions

In this specification, where it is not explicitly stated, all multi-byte values are stored with the most significant byte first. For example, in a 4 byte field, byte 0 will contain the MSB and byte 3 the LSB.

17.4 FORMAT UNIT (04)

Table 97: FORMAT UNIT (04)

Byte	BIT							
	7	6	5	4	3	2	1	0
0	Command Code = 04h							
1	FMTPINFO		LONG LIST=0	FMTDATA	CMPLIST	Defect List Format		
2	VU = 0							
3-4	Obsolete = 0							
5	VU = 0		Reserved = 0			FLAG	LINK	

- **FMTPINFO (Format Protection Information)** in combination with the Protection Field Usage field in the Parameter List Header specifies whether or not the drive enables or disables the use of protection information (see table defined in the Parameter List Header section).
- **FmtData** set to one specifies that a Data Out phase follows the Command phase. The Data Out phase consists of a Parameter List header, optionally followed by an Initialization Pattern Descriptor, optionally followed by a Defect List. If FmtData=0, the following defaults are assumed: DPROY=0, DCRT=1, STPF=1, IP=0, DSP=0, Immed=0.
- **CmpLst** is ignored.
- **Defect List Format** is ignored.
- **Notes:** It is recommended that the MODE SELECT command be issued prior to the FORMAT UNIT command to specify parameters that affect the formatting process.

The Block Length parameter of the Mode Select Parameter List's Block Descriptor is used during formatting and is saved following a successful format operation. If a MODE SELECT command has not been issued since the last reset or start-up (bring-up) sequence, then the Block Length from the previous format operation is used.

Subsequent to receiving a FORMAT UNIT command, the Target responds to commands as follows:

- All commands except REQUEST SENSE and INQUIRY return *Check Condition* status, while the format operation is an active I/O process.
- When tagged queuing is enabled (DQue = 0), all commands except REQUEST SENSE and INQUIRY return *Queue Full* status, while the FORMAT UNIT command is a queued I/O process.
- When tagged queuing is disabled (DQue = 1), all commands except REQUEST SENSE and INQUIRY return *Busy* status, while the FORMAT UNIT command is a queued I/O process.
- If a REQUEST SENSE command is received while a format operation is an active I/O process, the Target returns *Good* status. The sense key is set to *Not ready* and the additional sense code and qualifier is set to *Format In Progress*.
- If an INQUIRY command is received while a format operation is an active I/O process, the Target returns *Good* status and Inquiry data as requested.

The format operation must complete successfully for the Drive to be usable. If the command is interrupted by a reset, power down, or an unrecoverable error, the Drive enters a degraded mode of operation in which reading and writing are prohibited. To exit the degraded mode, another FORMAT UNIT command must be sent by the Initiator and completed successfully by the Target.

The FORMAT UNIT command sets the *Unit Attention Condition* for all Initiators except the one that issued the FORMAT UNIT command.

17.4.1 Parameter List Header

Following is the format of the Parameter List Header sent during the data out phase when FmtData is set to one.

Table 98: Format of the Parameter List Header

Byte	BIT							
	7	6	5	4	3	2	1	0
0	Reserved = 0				Protection Field Usage			
1	FOV	DPRY	DCRT	STPF = 1	IP	Obsolete	Immed	Vendor-specific
2 -3	(MSB) Defect List Length				(LSB)			
4-n	Initialization Pattern Descriptor							
(n+1) - m	Defect Descriptor							

- **Protection Field Usage:** in combination with the format protection information (FMTPINFO) field in the CDB specifies whether or not the drive enables or disables the use of protection information:

FMTPINFO	Protection Field Usage	Description
00h	000h	The drive will be formatted to type 0 protection
01h	xxxh	Check Condition status will be returned with the sense key set to Illegal Request and the additional sense code set to Invalid Field in the CDB.
10h	000h	The drive will be formatted to type 1 protection
11h	000h	The drive will be formatted to type 2 protection
11h	001h	Type 3 protection is not supported - Check Condition status will be returned with the sense key set to Illegal Request and the additional sense code set to Invalid Field in the Parameter List

All other combinations of FMTPINFO and Protection Field Usage will result in Check Condition status to be returned with the sense key set to Illegal Request and the additional sense code set to Invalid Field in the Parameter List.

Type 0 protection specifies that the drive shall disable the use of protection information and format to the block size specified. Following a successful format, the PROT_EN bit in the READ CAPACITY (16) parameter data will indicate that protection information is disabled.

Type 1 and type 2 protection specifies that the drive shall enable the use of protection information and format to the block size specified + 8 (e.g., if the block length is 512, then the formatted block length is 520). See format of data below. When protection information is written during a FORMAT UNIT command, protection information shall be written with a default value of all 0xFF's. Following a successful format, the PROT_EN bit in the READ CAPACITY (16) parameter data will indicate that protection information is enabled and the P_TYPE field in the READ CAPACITY (16) parameter data will indicate the protection type.

Byte	BIT							
	7	6	5	4	3	2	1	0
0... n	User Data							
n... n+1	Logical Block Guard							
n+2... n+3	Logical Block Application Tag							
n+4... n+7	Logical Block Reference Tag							

- **The Logical Block Guard field** contains a CRC that covers the preceding user data. This field is generated/checked per the SBC standard.
- **The Logical Block Application Tag field** may be modified by the initiator if the ATO bit is set to zero in mode page 0x0A. If the ATO bit is set to one, then the initiator shall not modify the Logical Block Application Tag field. This field is generated/checked per the SBC standard.

- **The Logical Block Reference Tag field** is generated/checked depending on protection types. With Type 1 protection, the Logical Block Reference Tag in the first logical block of the data transfer shall contain the least significant four bytes of the LBA contained in the Logical Block Address field of the command. Subsequent blocks shall contain the previous logical block reference tag plus one. With Type 2 protection, the Logical Block Reference Tag in the first logical block of the data transfer shall contain the value in the Expected Initial Logical Block Reference Tag field of the command. Subsequent blocks shall contain the previous logical block reference tag plus one.
- **FOV** (Format Options Valid) bit set to zero indicates that the Target should use its default settings for the DPRY (0), DCRT (1), STPF (1), IP (0), and DSP (1) bits. These bits must all be set to zero in the Parameter List Header when FOV=0, or the command will be terminated with Check Condition status, sense key of Illegal Request, and additional sense code of Invalid Field in Parameter List. FOV=1 indicates that the values set in DPRY, DCRT, STPF, IP, and DSP will be defined as specified below.
- **DPRY** (Disable Primary) bit set to zero disables error injection mode. A DPRY bit set to one enables error injection mode
- **DCRT** (Disable Certification) is ignored.
- **STPF** (Stop Format) is ignored.
- **IP** (Initialization Pattern) bit set to zero specifies that an initialization pattern descriptor is not included and all customer data will be initialized to zeroes. An IP bit of one specifies that an Initialization Pattern Descriptor is included in the FORMAT UNIT parameter list following the parameter list header. The Initialization Pattern Descriptor provides a means of enabling the Security Initialize option, which is not enabled by default. If anything in the Initialization Pattern Descriptor is not set as specified below, the command will be immediately terminated with Check Condition status, sense key of Illegal Request, and additional sense code of Invalid Field in Parameter List.

Table 99: Initialization Pattern Descriptor:

Byte	BIT							
	7	6	5	4	3	2	1	0
0	IP Modifier = 0		SI = 1	Reserved = 0				
1	IP Type = 0							
2-3	Initialization Pattern Length = 0							

-
- **IP Modifier** must be set to 0, indicating that the drive will not modify the initialization pattern.
- **SI (Security Initialize)** bit must be set to one when an Initialization Pattern Descriptor is sent. This specifies that the drive shall attempt to erase all locations that may contain customer data, including known defects.
- **Initialization Pattern Type** must be zero, indicating that the drive will use the default initialization pattern. All customer data will be initialized to zeroes.
- **Initialization Pattern Length** must be zero, as user-specified initialization patterns are not supported.
- **DSP** (Disable Saving Parameters) bit when zero indicates the target is to save all the current MODE SELECT saveable parameters during the format operation. When the bit is one, the target is not to save the current MODE SELECT saveable parameters.
- **Immed** (Immediate) bit set to zero requests that status be returned at the end of the format operation. An immediate bit set to one requests that status be returned immediately following CDB validation and transfer of data in the Data Out phase. If the format operation, with the immediate bit set to one, terminates in error, DEFERRED ERROR SENSE data is generated.
- **Defect List Length** must be zero. A user-supplied defect list is not supported. Otherwise the 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 Parameter List.

17.5 INQUIRY (12)

Table 100: INQUIRY (12)

Byte	BIT							
	7	6	5	4	3	2	1	0
0	Operation Code = 12h							
1	Reserved = 0						CmdDT =0	EVPD
2	Page Code							
3 - 4	Allocation Length							
5	VU = 0		Reserved = 0				FLAG	LINK

The INQUIRY command requests the parameters of the target to be sent to the Initiator.

An **EVPD** An EVPD bit of one specifies that the target return the vital product data page identified by the Page Code field in the CDB. The available VPD pages are defined in the addendum provided for each different drive model in the section entitled Inquiry Data Format.

The **Page Code** specifies which page of vital product data information the drive shall return.

Table 101: Page Code descriptions

EVPD	PAGE CODE	Description
0	0	The Target returns the standard INQUIRY data.
0	Non Zero	The drive returns <i>Check Condition</i> status with the sense key of <i>Illegal Request</i> and the additional sense code of <i>Invalid Field in CDB</i> .
1	Non Zero	The drive returns the vital product data of page code requested.

Allocation Length specifies the number of bytes that the Initiator has allocated for INQUIRY data to be returned. An allocation length of zero implies that no data is to be returned. The Target will terminate the DATA IN phase when all available INQUIRY data has been transferred or when allocation length bytes have been transferred, whichever is less.

Note: If an INQUIRY command is received from an Initiator with a pending unit attention condition (before the target reports *Check Condition* status), the Target processes the INQUIRY command. The unit attention condition is not cleared by this action.

Note: The INQUIRY command is a Priority command and is not queued.

Note: The inquiry data is set at the time of manufacture and will not change (without a FRU change), with the following exceptions:

- Product Revision Level (EVPD=0) can be changed when microcode is downloaded with the Write Buffer command.
- The information returned for EVPD=1, Page Code = 3 is not fixed.

Note: The inquiry data returned when media is not available will not be complete.

Byte 0 of the returned data on an INQUIRY command is the same no matter which page(s) is(are) returned. This description is to be used for all the following page definitions.

The Peripheral Qualifier field of zero (0) indicates that the peripheral device is currently connected to this logical unit. A Peripheral Device Type field of zero (0) indicates that this device is a Direct Access Storage Device (DASD).

17.5.1 Inquiry Data

Fields with a value shown inside quotes (e.g. Value = 'xyz') are character fields. A value not in quotes is a numeric value. Character fields are alphanumeric and represented in either ASCII.

17.5.1.1 Inquiry Data Format - EVPD = 0, Page Code = 0

Table 102: Inquiry Data- EVPD = 0

Byte	BIT							
	7	6	5	4	3	2	1	0
0	Qualifier = 0			Peripheral Device Type = 0				
1	RMB = 0	Reserved=0						
2	Version = 6							
3	Obsolete	Obsolete	Norm ACA=0	HiSup = 1	Response Data Format = 2			
4	Additional Length = 159 (9Fh)							
5	SCCS=0	ACC=0	TPGS=00b		3PC=0	Reserved = 0		Protect=1
6	Obsolete	EncSer = 1	Port	MultiP=1	Obsolete			RSVD = 0
7	Obsolete	Obsolete	RSVD = 0	RSVD = 0	Obsolete	Obsolete	CmdQue=1	RSVD = 0
8-15	Vendor ID = "HGST "(ASCII)							
16-31	Product ID (ASCII)							
32-35	Product Revision Level (ASCII)							
36-43	Unit Serial Number (ASCII)							
44-95	Reserved = 0							
96	Hard Assigned ALPA							
97	Acquired ALPA							
98-147	Copyright Notice (ASCII)							
148-163	Reserved=0							

- **Qualifier** is set to zero to indicate that the LUN specified is currently supported. Qualifier is set to 011b when the LUN specified is not present ¹
- **Peripheral Device Type** is set to zero to indicate that the device is a Direct-Access Peripheral Device.
- **Removal Media Bit (RMB)** is always set to zero to indicate no removal media exists.
- **Version** indicates the level of the ANSI standard that the product supports. The drive supports ANSI SPC-4.
- **NormACA** (Normal ACA) field of 0 indicates the device server does not support setting the NACA bit to one in the Control Byte of the CDB as defined in the SAM.
- **HiSup** bit of 1 indicates that the drive uses the hierarchical addressing model to assign LUNs to logical units.

1.If an INVALID LUN is specified, a *Check Condition* status will be returned for all commands except INQUIRY and REQUEST SENSE.

- **Response Data Format** is set to two to indicate that the INQUIRY Data Format as specified in the ANSI SCSI version 2 is supported by the Target.
- **Additional Length** indicates the number of bytes of INQUIRY information that follows.
- **SCCS** bit of zero indicates that the device does not contain an embedded storage array controller component.
- **ACC** bit of zero indicates that no access controls coordinator may be addressed through this logical unit.
- **TGPS** field of zero indicates that the device does not support asymmetric logical unit access.
- **3PC** bit of zero indicates that the device does not support third-party copy commands.
- **Protect** bit of one indicates that the drive supports protection information
- **EncSer** (Enclosure Services) bit of 0 indicates that the Target does not contain an embedded enclosure services component.
- **Port** bit of 0 indicates that the drive received the Inquiry command on port A, while a Port bit of 1 indicates that the drive received the Inquiry command on port B.
- **MultiP** (MultiPort) bit of 1 indicates that the Target has multiple ports and implements multi-port requirements.
- **CmdQue** is set to one to indicate that the drive supports command queuing.
- **Vendor ID** is HGST padded with ASCII blanks.
- **Product ID** is specified in 3 of Section 4.3.1.
- **Product Revision Level** indicates the level of microcode.
- **Unit Serial Number** contains the drive serial number.
- **Hard Assigned ALPA** contains the hard Pretreated Loop Physical Address of the port which received the Inquiry command.
- **Acquired ALPA** contains the Arbitrated Loop Physical Address acquired by the port which received the Inquiry command.

17.5.1.2 Inquiry Data Format - EVPD = 1 - Page Code = 00h

Table 103: Inquiry Data - EVPD = 1 (Page Code = 00h)

Byte	BIT							
	7	6	5	4	3	2	1	0
0	Qualifier = 0			Peripheral Device Type = 0				
1	Page Code = 00h							
2	Reserved = 0							
3	Page Length = 0Ah							
4	Supported Page Code = 00h							
5	Supported Page Code = 03h							
6	Supported Page Code = 80h							
7	Supported Page Code = 83h							
8	Supported Page Code = 86h							
9	Supported Page Code = 87h							
10	Supported Page Code = 88h							
11	Supported Page Code = 8Ah							
12	Supported Page Code = B1h							
13	Supported Page Code = D2h							

- **Qualifier** is set to zero to indicate that the LUN specified in the Command Block is currently supported.
- **Peripheral Device Type** is set to zero to indicate that the device is Direct Access.
- **Page Code** is set to 0, and this field contains the same value as in the page code field of the INQUIRY command descriptor block.
- **Page length** specifies the length of the following page data.
- **Supported Page Code** field contains the Page Codes supported by the Target. The list is in ascending order.

17.5.1.3 Inquiry Data Format - EVPD = 1, Page Code - 03h

Table 104: Inquiry Data - EVPD = 1 (Page Code = 03h)

Byte	BIT							
	7	6	5	4	3	2	1	0
0	Qualifier = 0			Peripheral Device Type = 0				
1	Page Code = 03h							
2	Reserved = 0							
3	Page Length = 204 (CCh)							
4	ASCII Fields Length = 00h							
5-7	Reserved = 0							
8-23	Reserved = 0							
24-35	ASCII uCode Identifier							
36-39	Reserved = 0							
40-41	Major Version							
42-43	Minor Version							
44-47	User Count							
48-51	Build Number							
52-79	Build Date String							
80-81	Code ID							
82-83	Compatibility ID							
84-91	Product ID							
92-99	Interface ID							
100-107	Code Type							
108-119	User Name							
120-135	Machine Name							
136-167	Directory Name							
168-171	Operating State							
172-175	Functional Mode							
176-179	Degraded Reason							
180-183	Broken Reason							
184-187	Code Mode							
188-195	ASCII uCode revision							
196-199	Context failure reason							
200-203	South Assert Address							
204-205	North Assert Code							
206-207	Reserved							

- **Qualifier** is set to zero to indicate that the LUN specified in the Command Block is currently supported.
- **Peripheral Device Type** is set to zero to indicate that the device is Direct Access.
- **Page Code** is set to the value of the page code field in the CDB.
- **Page Length** field specifies the length (in bytes) of the vendor unique VPD information (bytes 4 - 163). If the allocation length of the CDB is too small to transfer all the data, the Page Length field is not adjusted to reflect the truncation.
- **ASCII uCode Identifier** contains the drive's microcode identifier. The field is alphanumeric (ASCII), left aligned, and the unused bytes are ASCII spaces (20h).
- **Major Version** and **Minor Version** are version numbers of the code loaded on the drive.
- **User Count** is the number of times the code has been built since the master build.

- **Build Number** is the master build version number.
- **Build Date String** is the date the code on the drive was built, in an extended string format.
- **Code ID** is a binary value for firmware development tracking.
- **Compatibility ID** is a binary value for firmware development tracking.
- **Product ID** is the name of the product this code is for.
- **Interface ID** is the interface type and serial interface speed (e.g. SCSI or FCAL 4Gb) of the code.
- **Code Type** is the intended use of the this code. (e.g. local, released, test)
- **User Name** is the username of the person who built this version of the code.
- **Machine Name** is the workstation on which this version of the code was built.
- **Directory Name** is the last 32 characters of the directory from where this code was built.
- **Operating State** is the drive operating state. The least significant bit contains the following:

0 = OM_BROKEN	We have detected a hardware failure or there was an error loading context.
1 = OM_DEGRADED	We have a soft failure; i.e., incomplete format.
2 = OM_INACCESSIBLE	Drive is good.
3 = OM_STARTING	Loading context..
5 = OM_NORMAL	Context is loaded and ready to read/write.
7 = OM_STOPPED	Drive has come ready but now has been stopped.
- **Functional Mode** is the drive functional mode. The least significant byte (0x0000000n) contains the following:

0 = OM_NORMAL_MODE	Not in special or recovery mode.
1 = OM_SPECIAL_CMD	Special command mode on.
3 = OM_SPC_RSV_ACCESS	Special cmd mode and access to reserved area allowed.
5 = OM_SPC_SDWNL0AD	Special cmd mode and special download allowed.
7 = OM_SPC_RACCESS_SDWNLD	Special cmd, access to reserved area, and special download allowed.

 The second byte (0x000n0000) contains the following:

0 = Idle functions are not enabled.	
1 = Idle functions are enabled.	
- **Degraded Reason** (UECType) is why the file is in a degraded mode; i.e., how to exit this mode.
- **Broken Reason** (UECType) is why the drive believes the hardware is broken.
- **Code Mode** is the type of code the drive is running. The least significant bit contains the following:

- 0 = OM_FLASH	Drive is running flash code
- 1 = OM_FLASH_OVERLAY	Drive is running flash overlay code
- 2 = OM_DISK	Drive is running code that has been loaded from NAND
- 3 = OM_TRANSIENT	Drive is running code that has been downloaded but not saved.
- **ASCII uCode revision** is the revision level of the media access firmware. This field is alphanumeric.
- **Context failure reason** is a qualifier when the broken reason indicates a damaged context.
- **South Assert Address** is used for debug of firmware asserts.
- **North Assert Code** is used for debug of firmware asserts.

17.5.1.4 Inquiry Data Format - EVPD = 1, Page Code - 80h

Table 105: Inquiry Data - EVPD = 1 (Page Code = 80h)

Byte	BIT							
	7	6	5	4	3	2	1	0
0	Qualifier = 0			Peripheral Device Type = 0				
1	Page Code = 80h							
2	Reserved = 0							
3	Page Length = 16 (10h)							
4-19	Serial Number (ASCII)							

- **Qualifier** is set to zero to indicate that the LUN specified in the Command Block is currently supported.
- **Peripheral Device Type** is set to zero to indicate that the device is Direct Access.
- **Page Code** is set to the value of the page code field in the CDB.
- **Page Length** is set to 16, and this field specifies the length of the following page data.
- **Serial Number** gives the drive serial number, right aligned.

17.5.1.5 Inquiry Data - EVPD = 1 (Page Code = 83h)

Table 106: Inquiry Data - EVPD = 1 (Page Code = 83h)

Byte	BIT							
	7	6	5	4	3	2	1	0
0	Qualifier = 0			Peripheral Device Type = 0				
1	Page Code = 83h							
2	Reserved = 0							
3	Page Length = 12 (0Ch)							
4	Reserved = 0				Code Set = 1			
5	Reserved = 0		Association = 0		Identifier Type = 3			
6	Reserved = 0							
7	Identifier Length = 8							
8-15	(MSB) World Wide ID (LSB)							

- **Qualifier** is set to zero to indicate that the LUN specified in the Command Block is currently supported.
- **Peripheral Device Type** is set to zero to indicate that the device is Direct Access.
- **Page Code** is set to the value of the page code field in the CDB.
- **Page Length** is set to 12, and this field specifies the length of the following page data.
- **Code Set** field specifies the code set used for the identifier field. The Target supports binary.
- **Association** field is set to 0, indicating that the Identifier field is associated with the logical unit.
- **Identifier Type** field specifies the format and assignment authority for the identifier. The Target supports the value of 03h.
- **World Wide ID** is a 64-bit unique value for each drive. The format is: **5000CCAh xxxh nb yyb** where:
 - xxx** is the 12-bit Block Assignment defined for each model and manufacturing site
 - n** is the 22-bit drive unique serial number representation
 - yy** is the 2-bit Port Identifier

17.5.1.6 Inquiry Data Format - EVPD = 1, Page Code - 86h

Table 107: Inquiry Data Format - EVPD = 1, (Page Code - 86h)

Byte	BIT							
	7	6	5	4	3	2	1	0
0	Qualifier = 0			Peripheral Device Type = 0				
1	Page Code = 86h							
2	Reserved = 0							
3	Page Length = 60 (3Ch)							
4	Reserved = 0		SPT			GRD_CHK	APP_CHK	REF_CHK
5	Reserved = 0			Group_Sup	Prior_Sup	HEADSUP	ORDSUP	SIMPSUP
6	Reserved=0						NV_SUP	V_SUP
7-63	Reserved = 0							

- **SPT (Supported Protection Type)** field is set to 001b to indicate that the drive supports type 1 and type 2 protection.
- **GRD_CHK (Guard Check)** is set to 1 to indicate that the drive checks the Logical Block Guard Tag field in the protection information, if any.
- **APP_CHK (Application Tag Check)** bit is set to 1 to indicate that the drive checks the Logical Block Application Tag field in the protection information, if any.
- **REF_CHK (Reference Tag Check)** bit is set to 1 to indicate that the drive checks the Logical Block Reference Tag field in the protection information, if any.
- **GROUP_SUP (Group Supported)** bit is set to 0 to indicate that the grouping function is not supported.
- **PRIOR_SUP (Priority Supported)** bit is set to 0 to indicate that task priority is not supported.
- **HEADSUP (Head of Queue Supported)** bit is set to 0 to indicate that Head of Queue is not supported.
- **ORDSUP (Ordered Supported)** bit is set to 0 to indicate that Ordered task is not supported.
- **SIMPSUP (Simple Supported)** is set to 1 to indicate support for Simple task attributes.
- **NV_SUP (Non-volatile Supported)** is set to 0 to indicated that non-volatile cache features are not supported.
- **V_SUP (Volatile Supported)** is set to 1 to indicated support of a volatile cache.

17.5.1.7 Inquiry Data Format - EVPD = 1, Page Code - 87h

Byte	BIT							
	7	6	5	4	3	2	1	0
0	Qualifier = 0			Peripheral Device Type = 0				
1	Page Code = 87h							
2-3	Page Length = 0004h							
4	Reserved=0		Policy Page Code = 3Fh					
5	Policy Subpage Code = FFh							
6	MLUS=1	Reserved = 0					Mode PagePolicy = 0	
7	Reserved = 0							

- **Policy Page Code** set to 3Fh and **Policy Subpage Code** set to FFh indicate that the descriptor applies to all mode pages and subpages
- **MLUS (Multiple Logical Units Share)** set to 1 indicates the policy is shared by multiple logical units.

- **Mode Page Policy** set to 00b indicates that all mode pages and subpages are shared.

17.5.1.8 Inquiry Data Format - EVPD = 1, Page Code - 88h

Byte	BIT							
	7	6	5	4	3	2	1	0
0	Qualifier = 0			Peripheral Device Type = 0				
1	Page Code = 88h							
2-3	Page Length = 48 (0030h)							
4-5	Reserved=0							
6-7	Primary Relative Port = 0001h							
8-9	Reserved = 0							
10-11	Initiator Port Transport ID Length = 0							
12-13	Reserved = 0							
14-15	Primary Target Port Descriptors Length = 0Ch							
16	Protocol Identifier				Code Set = 1			
17	PIV=1	RSVD	Association = 1		Identifier Type = 3			
18	Reserved = 0							
19	Identifier Length = 8							
20-27	(MSB) Primary Target Port Identifier (World Wide ID)							(LSB)
28-29	Reserved = 0							
30-31	Secondary Relative Port = 0002h							
32-33	Reserved = 0							
34-35	Initiator Port Transport ID Length = 0							
36-37	Reserved = 0							
38-39	Secondary Target Port Descriptors Length = 0Ch							
40	Protocol Identifier				Code Set = 1			
41	PIV=1	RSVD	Association = 1		Identifier Type = 3			
42	Reserved = 0							
43	Identifier Length = 8							
44-51	(MSB) Secondary Target Port Identifier (World Wide ID)							(LSB)

- **Protocol Identifier** is valid only when PIV=1. Protocol Identifier = 0 indicates Fibre Channel devices. Protocol Identifier = 6 indicates SAS devices
- **Code Set** specifies the data type for the identifier field. Code Set = 1 indicates binary data
- **PIV (Protocol Identifier Valid)** set to one indicates that the Protocol Identifier field contains a valid value.
- **Association** specifies the entity with which the Identifier field is associated: 1h for Target or Relative Port.
- **Identifier Type** specifies the format and assignment authority for the identifier: 3h indicates NAA format of the WWID for Target Port.
- **Identifier** fields contain the actual Identifier Descriptor.
 - The Target Port Identifiers are defined in the NAA IEEE WWID format where:
 World Wide ID is a 64-bit unique identification for each drive. The format is: 5000CCAh xxxh yyb n where xxx is the 12-bit block assignment defined for each model and manufacturing site yy is the 2-bit port/node ID select n is the 22-bit drive unique serial number.

17.5.1.9 Inquiry Data Format - EVPD = 1, Page Code - B1h

Table 108: Inquiry Data - EVPD = 1 (Page Code = B1h)

Byte	BIT							
	7	6	5	4	3	2	1	0
0	Qualifier = 0			Peripheral Device Type = 0				
1	Page Code = B1h							
2-3	Page Length = 60 (003Ch)							
4-5	Medium Rotation Rate = 1 (1h)							
6	Reserved							
7	Reserved				Nominal Form Factor = 3h			
8-63	Reserved							

- **Qualifier** field is set to zero to indicate that the LUN specified in the Command Block is currently supported.
- **Peripheral Device Type** is set to zero to indicate that the device is Direct Access.
- **Page Code** is set to the value of the page code field in the CDB.
- **Page Length** is set to 60, and this field specifies the length of the following page data.
- **Medium Rotation Rate** field is set to 1, which indicates the drive is an SSD.
- **Nominal Form Factor** field is set to 3h.

17.5.1.10 Inquiry Data Format - EVPD = 1, Page Code - D2h

Table 109: Inquiry Data - EVPD = 1 (Page Code = D2h)

Byte	BIT							
	7	6	5	4	3	2	1	0
0	Qualifier = 0			Peripheral Device Type = 0				
1	Page Code = D2h							
2	Reserved = 0							
3	Page Length = 120 (78h)							
4	HDC Version Length = 19 (13h)							
5 - 23	ASCII HDC Version							
24	Card Serial Number Length = 19 (13h)							
25 - 43	ASCII Card Serial Number							
44	NAND FLASH Version Length = 19 (13h)							
45 - 63	ASCII NAND FLASH Version							
64	Card Assembly Part Number Length = 19 (13h)							
65 - 83	ASCII Card Assembly Part Number							
84	Second Card Serial Number Length = 19 (13h)							
85 - 103	ASCII Second Card Serial Number							
104	Second Card Assembly Part Number Length= 19 (13h)							
105 - 123	ASCII Second Card Assembly Part Number							

- **Qualifier** is set to zero to indicate that the LUN specified in the Command Block is currently supported.
- **Peripheral Device Type** is set to zero to indicate that the device is Direct Access.
- **Page Code** is set to the value of the page code field in the CDB.
- **Page Length** is set to 100, and this field specifies the length of the following page data.

Note: If the media is not available, bytes 0 through 3 are valid. All the other fields are ASCII blanks (20h) with a null terminator (00h).

Note: All ASCII fields are alphanumeric, left aligned, and padded on the right with ASCII blanks (20h) with a null terminator (00h).

17.6 LOG SELECT (4C)

Table 110: Log Select (4C)

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Command Code = 4Ch							
1	Reserved = 0			Reserved = 0			PCR	SP
2	PC		Page Code					
3	SubPage Code = 0							
4-6	Reserved = 0							
7	(MSB) Parameter List Length = 0 (LSB)							
8								
9	Reserved = 0						FLAG	LINK

The LOG SELECT command provides a means for the Initiator to clear statistical information maintained by the drive and reported via the LOG SENSE command.

- **PCR** The Parameter Code Reset determines whether the Log Sense parameters will be cleared and unit attention posted for all other Initiators. A value of 1 indicates that the parameters be cleared, while a value of zero (except when PC = 11b) indicates that the parameters not be cleared. Parameter list length must be zero when PCR is 1. The PC field is ignored for list parameters, i.e. when the Format and Linking (F&L) field contains 01b or 11b.
- **SP** The Save Parameters bit value of zero indicates that the page parameters not be saved. A value of 1 indicates that the page parameters that are savable be saved after they have been changed. SP bit MUST be 1 if parameter list length is greater than zero. Otherwise it will result in a *Check Condition* status being returned. The sense key shall be set to *Illegal Request* and additional sense code of *Invalid Field in CDB*.
- **PC** The Page Control field defines the type of parameters to be selected. The PC field set to 11b (and PCR is then a don't care) indicates that the Default Cumulative values are set to their default values of 0. If the PC field is set to 01b and PCR is set to 1, the Current Cumulative values are also set to their default values of 0.

Parameter List Length MUST be zero when PC = 11b. Otherwise the command is terminated and a *Check Condition* status is returned. The sense key shall be set to *Illegal Request* and additional sense code of *Invalid Field in CDB*.

- **Page Code** field identifies which page is being selected. This field must be set to the values indicated in Page 0. If the Page Code value is invalid a Check Condition status is returned with a sense key of Illegal Request and additional sense code of Invalid Field in CDB.
 - If page code field is set to zero, then the selection applies to all log parameters in all valid log pages.
 - If page code field is set to a non zero, then the selection applies to all log parameters specified by this field.
- **SubPage Code** This field specifies the subpage to select. This field is not supported and must be set to 0.
- **Parameter List Length** The Parameter List Length field specifies the length in bytes of the parameter list that shall be located in the DATA OUT buffer. A parameter list length zero indicates that no pages shall be transferred.
 - If the PARAMETER LIST LENGTH field is set to zero, then the PCR bit, the SP bit, and the PC fields apply to the page (pages) addressed by the page code field.
 - If The PARAMETER LIST LENGTH field is set to non zero, and the if PAGE CODE field is non-zero or the SUBPAGE CODE field is non-zero, then the command shall be terminated with CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST, and the additional sense code set to INVALID FIELD IN CDB.

•

Note: A specified length greater than 0x00FF will result in a *Check Condition* status being returned. A length that results in log data being truncated will generate a *Check Condition* status.

Note: For page 0Fh, the maximum parameter list length supported is 4004h (4 bytes for the header and 100h bytes for each of the 40h parameters that are supported). The Parameter List Length must be an integral of the number of parameters plus the 4 byte header. (Ex: Parameter length =104h for one parameter, 204h for 2 parameters,... 4004h for all 40h parameters).

The drive allows updates to the current cumulative values only. A value of zero is accept and is not considered an error. The drive updates only pages 0Eh, the Start/Stop Cycle page and 0Fh, the Application Client page. For other pages the parameters are ignored. If the data out buffer contains multiple pages then the application client should send the pages in ascending order. If the data out buffer contains multiple log parameters within a page, all log parameters within the page should be sent and they should be sent in ascending order by parameter code value. The drive shall return Check Condition status if the application client sends pages out of order, parameter codes out of order or missing parameter code. The sense key shall be set to Illegal Request and additional sense code set to Invalid Field in Parameter List. If one or more fields of the CDB are not set correctly the command will be terminated with a *Check Condition* status. The sense key shall be set to *Illegal Request* and additional sense code of *Invalid Field in CDB*. To indicate that parameters have changed, the Target generates a unit attention condition for all Initiators except the one that issued the LOG SELECT command.

The following list contains all individual page parameters (counters) that are set to their default value of zero by the LOG SELECT command (when PCR=1).

- Page **02h** parameters: (Counters for write errors)
 - Write errors recovered without delay
 - Write errors recovered with possible delays
 - LBAs with write fault error
 - Total errors recovered
 - Number of times recovery invoked
 - Total write byte count
 - LBAs with hard error
- Page **03h** parameters: (Counters for read errors)
 - Read errors recovered without delay
 - Read errors recovered with possible delays
 - LBAs with ECC detected error
 - Total errors recovered
 - Number of times recovery invoked
 - Total read byte count
 - LBAs with hard error.
- Page **05h** parameters: (Counters for Verify Errors)
 - Errors recovered without delay
 - Errors recovered with possible delays
 - LBAs with ECC detected error
 - Total errors recovered
 - Number of times recovery invoked
 - Total bytes verified

- LBAs with hard error.
- Page **06h** parameters: (Counters for non medium errors, seek and other hardware type failures)
 - Non-Medium Error Counter
- Page **15h** parameters: (Background Medium Scan information)
 - BMS Status parameter
 - all Medium Scan parameters
- Page **30h** parameters:
 - Overrun Counter
 - Under run Counter
 - Device Cache Full Read Hits
 - Device Cache Partial Read Hits
 - Device Cache Write Hits
 - Device Cache Fast Writes
 - Device Cache Misses on Reads
- Page **37h** parameters:
 - Media Exception
 - Hardware Exception
 - Total Read Commands
 - Total Write Commands

17.7 LOG SENSE (4D)

Table 111: Log Sense (4D)

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Command Code = 4Dh							
1	Reserved = 0			Reserved = 0			PPC=0	SP
2	PC	Page Code						
3	Subpage Code = 0							
4	Reserved = 0							
5	(MSB) Parameter Pointer = 0 (LSB)							
6								
7	(MSB) Allocation Length (LSB)							
8								
9	Reserved = 0					FLAG	LINK	

The LOG SENSE command allows the Initiator to retrieve the statistical data regarding the drive.

- **PPC** (Parameter Pointer Control) bit must be set to zero. This specifies that the drive start transferring data starting from the field specified in the parameter pointer field for the number of bytes specified by the allocation length. If the PPC bit is set to 1, *Check Condition* status is returned with a sense key of *Illegal Request* and additional sense code of *Invalid Field in CDB*.
- **SP** (Save Parameters) bit set to 0 specifies that the drive does not save any log parameters. If it is set to 1, all page parameters that are savable (those pages denoted by a DS = 0 in the parameter header control byte) are saved.
- **PC** (Page Control) field defines the type of parameters to be selected. This field must be set to 01b to specify the current cumulative values. Any other value in this field will cause the command to end with a *Check Condition* status with a sense key of *Illegal Request* and an additional sense code of *Invalid Field in CDB*.
- **Page Code** field identifies which page is being requested. This field must be set to the values indicated in Page 0. If the Page Code value is invalid a *Check Condition* status is returned with a sense key of *Illegal Request* and additional sense code of *Invalid Field in CDB*.
- **SubPage Code** This field specifies the subpage to select. This field is not supported and must be set to 0.
- **Parameter Pointer Field** specifies the beginning field for the transfer. This field must be set to 0000h. If the Parameter Pointer Field is not zero a *Check Condition* status is returned with a sense key of *Illegal Request* and additional sense code of *Invalid Field in CDB*.
- **Allocation Length** field specifies the maximum number of bytes the Initiator has allocated for returned Log Sense Data. No bytes are transferred if the length is zero. This condition is not considered an error. The Target terminates the Data In phase when all available Log Sense data has been transferred or when the number of bytes equals the allocation length, whichever is less.

17.7.1 Log Page parameters

Each log page begins with a 4-byte page header followed by zero or more variable-length log parameters.

Page header

Page Code field identifies which log page is being transferred.

The Page Length field specifies the length in bytes of the following log parameters.

Log parameters

Each log parameter begins with a 4-byte parameter header followed by one or more bytes of parameter value data.

The Parameter Code field identifies which log parameter is being transferred for that log page.

The Parameter Control field, the 3rd byte of each parameter header, contains several fields.

- **DU** The Disable Update bit is set to 0 to indicate that the drive updates the log parameter value to reflect events that should be noted by that parameter.
- **TSD** The Target Save Disable bit is set to zero to indicate that the drive provides a Target defined method for saving log parameters.
- **ETC** The Enable Threshold Comparison bit is set to 0 to indicate the drive does not perform comparisons between cumulative and any threshold values.
- **TMC** The Threshold Met Criteria field is not valid because this drive does not perform threshold comparisons. This field is set to 0.
- **Format and Linking** The F & L field indicates the type of log parameter and how parameters that reach their maximum value are handled.
 - 00b: Data counter: If any other parameter in this log page reaches its maximum value, then this parameter shall stop incrementing until reinitialized by a Log Select command.
 - 01b: List format ASCII data: No maximum values to handle
 - 10b: Data counter: If another parameter reported in this log page reaches its maximum value, then this parameter shall not stop incrementing. This parameter may be reinitialized by a Log Select command.
 - 11b: List format binary data: No maximum values to handle.

17.7.2 Log Sense Page 0

Page 0 indicates the supported log sense pages. This page is used to determine which additional pages an Initiator can request.

Table 112: Log Sense Page 0

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Reserved		Page code = 0					
1	Reserved							
2-3	Page Length = 000Fh (Number of Pages Supported)							
4	First supported page 00h							
5	Second supported page 02h							
6	Third supported page 03h							
7	Fourth supported page 05h							
8	Fifth supported page 06h							
9	Sixth supported page 0Dh							
10	Seventh supported page 0Eh							
11	Eighth supported page 0Fh							
12	Ninth supported page 10h							
13	Tenth supported page 11h							
14	Eleventh supported page 15h							
15	Twelfth supported page 1Ah							
16	Thirteenth supported page 2Fh							
17	Fourteenth supported Page Code =30h							
18	Fifteenth supported Page Code = 38h							

17.7.3 Log Sense Page 2

This page contains counters for write errors.

Table 113: Log Sense Page 2 (part 1 of 2)

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Reserved		Page code = 02h					
1	Reserved							
2-3	PageLength = 54h							
4-5	Parameter Code = 0000h							
6	DU = 0	DS = 0	TSD = 0	ETC = 0	TMC = 0		F&L = 00b	
7	Parameter Length = 08h							
8-15	Errors recovered without delay							
16-17	Parameter Code = 0001h							
18	DU = 0	DS = 0	TSD = 0	ETC = 0	TMC = 0		F&L = 00b	
19	Parameter Length = 08h							
20-27	Errors recovered with possible delays							
28-29	Parameter Code = 0002h							
30	DU = 0	DS = 0	TSD = 0	ETC = 0	TMC = 0		F&L = 00b	
31	Parameter Length = 08h							
32-39	Reserved = 0							
40-41	Parameter Code = 0003h							
42	DU = 0	DS = 0	TSD = 0	ETC = 0	TMC = 0		F&L = 00b	
43	Parameter Length = 08h							
44-51	Total errors recovered							
52-53	Parameter Code = 0004h							
54	DU = 0	DS = 0	TSD = 0	ETC = 0	TMC = 0		F&L = 00b	
55	Parameter Length = 08h							
56-63	Times recovery invoked							

Table 114: Log Sense Page 2 (part 2 of 2)

Byte	Bit							
	7	6	5	4	3	2	1	0
64-65	Parameter Code = 0005h							

66	DU = 0	DS = 0	TSD = 0	ETC = 0	TMC = 0	F&L = 00b
67	Parameter Length = 08h					
68-75	Total bytes written					
76-77	Parameter Code = 0006h					
78	DU = 0	DS = 0	TSD = 0	ETC = 0	TMC = 0	F&L = 00b
79	Parameter Length = 08h					
80-87	Count of hard errors					

All parameter counts indicate the number of sectors with the specified types of errors, except Times Recovery Invoked, which is a cumulative count of all recovery steps attempted on all sectors written.

17.7.4 Log Sense Page 3

This page contains counters for read errors.

Table 115: Log Sense Page 3 (part 1 of 2)

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Reserved		Page code = 03h					
1	Reserved							
2-3	PageLength = 54h							
4-5	Parameter Code = 0000h							
6	DU = 0	DS = 0	TSD=0	ETC = 0	TMC = 0		F&L = 00b	
7	Parameter Length = 08h							
8-15	Errors recovered without delay							
16-17	Parameter Code = 0001h							
18	DU = 0	DS = 0	TSD = 0	ETC = 0	TMC = 0		F&L = 00b	
19	Parameter Length = 08h							
20-27	Errors recovered with possible delays							
28-29	Parameter Code = 0002h							
30	DU = 0	DS = 0	TSD=0	ETC = 0	TMC = 0		F&L = 00b	
31	Parameter Length = 08h							
32-39	Reserved = 0							
40-41	Parameter Code = 0003h							
42	DU = 0	DS = 0	TSD=0	ETC = 0	TMC = 0		F&L = 00b	
43	Parameter Length = 08h							
44-51	Total errors recovered							
52-53	Parameter Code = 0004h							
54	DU = 0	DS = 0	TSD=0	ETC = 0	TMC = 0		F&L = 00b	
55	Parameter Length = 08h							
56-63	Times recovery invoked							
64-65	Parameter Code = 0005h							

Table 116: Log Sense Page 3 (part 2 of 2)

Byte	Bit							
	7	6	5	4	3	2	1	0

66	DU = 0	DS = 0	TSD = 0	ETC = 0	TMC = 0	F&L = 00b
67	Parameter Length = 08h					
68-75	Total bytes read					
76-77	Parameter Code = 0006h					
78	DU = 0	DS = 0	TSD = 0	ETC = 0	TMC = 0	F&L = 00b
79	Parameter Length = 08h					
80-87	Count of hard errors					

All parameter counts indicate the number of sectors with the specified types of errors, except Times Recovery Invoked, which is a cumulative count of all recovery steps attempted on all sectors read. ECC-on-the-fly correction is not included in any counters.

17.7.5 Log Sense Page 5

This page contains counters for verify errors.

Table 117: Log Sense Page 5 (part 1 of 2)

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Reserved		Page code = 05h					
1	Reserved							
2-3	PageLength = 54h							
4-5	Parameter Code = 0000h							
6	DU = 0	DS = 0	TSD = 0	ETC = 0	TMC = 0		F&L = 00b	
7	Parameter Length = 08h							
8-15	Errors recovered without delay							
16-17	Parameter Code = 0001h							
18	DU = 0	DS = 0	TSD = 0	ETC = 0	TMC = 0		F&L = 00b	
19	Parameter Length = 08h							
20-27	Errors recovered with possible delays							
28-29	Parameter Code = 0002h							
30	DU = 0	DS = 0	TSD = 0	ETC = 0	TMC = 0		F&L = 00b	
31	Parameter Length = 08h							
32-39	Reserved = 0							
40-41	Parameter Code = 0003h							
42	DU = 0	DS = 0	TSD = 0	ETC = 0	TMC = 0		F&L = 00b	
43	Parameter Length = 08h							
44-51	Total errors recovered							
52-53	Parameter Code = 0004h							
54	DU = 0	DS = 0	TSD = 0	ETC = 0	TMC = 0		F&L = 00b	
55.	Parameter Length = 08h							
56-63	Times recovery invoked							
64-65	Parameter Code = 0005h							

Table 118: Log Sense Page 5 (part 2 of 2)

Byte	Bit							
	7	6	5	4	3	2	1	0
66	DU = 0	DS = 0	TSD = 0	ETC = 0	TMC = 0		F&L = 00b	
67	Parameter Length = 08h							
68-75	Total Bytes Verified							
76-77	Parameter Code = 0006h							
78	DU = 0	DS = 0		TSD = 0	TMC = 0		F&L = 00b	
79	Parameter Length = 08h							
80-87	Count of hard errors							

All parameter counts indicate the number of sectors with the specified types of errors, except Times Recovery Invoked, which is a cumulative count of all recovery steps attempted on all sectors verified. ECC-on-the-fly correction is not included in any counters.

17.7.6 Log Sense Page 6

This page contains counters for non-medium errors. This includes seek errors and other hardware type failures.

Table 119: Log Sense Page 6

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Reserved		Page code = 06h					
1	Reserved							
2-3	PageLength = 0Ch							
4-5	Parameter Code = 00h							
6	DU = 0	DS = 0	TSD = 0	ETC = 0	TMC = 0		F&L = 00b	
7	Parameter Length = 08h							
8-15	Error count							

17.7.7 Log Sense Page D

This page contains temperature information.

Table 120: Log Sense Page D

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Reserved		Page code = 0Dh					
1	Reserved							
2-3	PageLength = 0Ch							
4-5	Parameter Code = 0000h							
6	DU = 0	DS = 1	TSD = 0	ETC = 0	TMC = 0		F&L = 00b	
7	Parameter Length = 02h							
8	Reserved							
9	Temperature (degrees Celsius)							
10-11	Parameter Code 0001h							
12	DU = 0	DS = 1	TSD = 0	ETC = 0	TMC = 0		F&L = 00b	
13	Parameter Length = 02h							
14	Reserved							
15	Reference Temperature (degrees Celsius)							

17.7.8 Log Sense Page E

This page contains manufacturing date information.

Table 121: Log Sense Page E

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Reserved		Page code = 0Eh					
1	Reserved = 0							
2-3	PageLength = 34h							
4-5	Parameter Code = 0001h							
6	DU=0	DS=1	TSD=0	ETC=0	TMC = 0		F&L = 00b	
7	Parameter Length = 06h							
8-11	Year of Manufacture (4 ASCII characters)							
12-13	Week of Manufacture (2 ASCII characters)							
14-15	Parameter Code 0002h							
16	DU=0	DS=0	TSD=0	ETC=0	TMC = 0		F&L = 00b	
17	Parameter Length = 06h							
18-21	Accounting Date Year (4 ASCII characters)							
22-23	Accounting Date Week (2 ASCII characters)							
24-25	Parameter Code 0003h							
26	DU=0	DS=1	TSD=0	ETC=0	TMC = 0		F&L = 00b	
27	Parameter Length = 04h							
28-31	Reserved = 0							
32-33	Parameter Code 0004h							
34	DU=0	DS=1	TSD=0	ETC=0	TMC = 0		F&L = 00b	
35	Parameter Length = 04h							
36-39	Reserved = 0							
40-41	Parameter Code 0005h							
42	DU=0	DS=1	TSD=0	ETC=0	TMC = 0		F&L = 00b	
43	Parameter Length = 04h							
44-47	Reserved = 0							
48-49	Parameter Code 0006h							

50	DU=0	DS=1	TSD=0	ETC=0	TMC = 0	F&L = 00b
51	Parameter Length = 04h					
52-55	Reserved = 0					

The week and year that the device was manufactured shall be set in the parameter field defined by parameter code 0001h. The date of manufacture cannot be saved using the LOG SELECT command. The data is expected in numeric ASCII characters (30-39h) in the form YYYYWW. The accounting date specified by parameter code 0002h is a parameter that can be saved using the LOG SELECT command.

17.7.9 Log Sense Page F

This page contains the Application Client Log.

Table 122: Log Sense Page F

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Reserved		Page code = 0Fh					
1	Reserved							
2-3	Page length = 4000h							
	Application client log parameter							
4-259	1st application client log parameter							
16132-16387	64th application client log parameter							

The following describes the application client log parameter structure.

Table 123: Log Sense Page F, Application Client Log

Byte	Bit							
	7	6	5	4	3	2	1	0
0-1	Parameter code							
2	DU = 1	DS = 0	TSD = 0	ETC = 0	TMC = 0		F&L = 00b	
3	Parameter length = FCh							
4-	First parameter byte							
255	Last parameter byte							

Parameter code 0000h through 003Fh are supported.

The values stored in the parameter bytes represent data sent to the device in a previous LOG SELECT command.

17.7.10 Log Sense Page 10

This page contains self-test results. The results of the 20 most recent self-tests are stored in this Log page.

Table 124: Log Sense Page 10

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Reserved		Page code = 10h					
1	Reserved							
2-3	PageLength = 190h							
4-23	1st self-test results log parameter							
384- 403	20th self-test results log parameter							

The following describes the self-test results log parameter structure

Table 125: Log Sense Page 10, self-test results

Byte	Bit								
	7	6	5	4	3	2	1	0	
0-1	Parameter code								
2	DU = 0	DS = 0	TSD = 0	ETC = 0	TMC = 0		F&L = 11b		
3	Parameter Length = 10h								
4	Function Code			RSVD	Self-Test Results Value				
5	Extended Segment Number								
6-7	Timestamp								
8-15	LBA of First Failure								
16	Reserved				Sense Key				
17	Additional Sense Code								
18	Additional Sense Code Qualifier								
19	Vendor specific								

- **Parameter Code** identifies the log parameter for the log page. The parameter code field for the results of the most recent test will be 0001h. The parameter for the next most recent will be 0002h.
- **Function Code** contains the content of the Function Code field in the SEND DIAGNOSTIC command that initiated this self-test.
- **Self-Test Results Value** is described in the table below.

Table 126: Log Sense Page 10, self-test results

Value	Description
0h	The self-test routine completed without error.
1h	The background self-test routine was aborted by the initiator using a SEND DIAGNOSTIC command with the Abort Background self-test function.
2h	The self-test routine was aborted by the application client by a Task Management function or a reset.
3h	An unknown error occurred while the Target was executing the self-test routine and the Target was unable to complete the self-test routine.
4h	The self-test completed with a test element that failed and it is not known which test element failed.
5h	The first segment of the self-test failed.
6h	The second segment of the self-test failed.
7h	The third or greater segment of the self-test failed (see the Extended segment number field).
8h-Eh	Reserved.
Fh	The self-test is in progress.

- **Extended Segment Number** This field identifies the number of the segment that failed during self-test. If no segment failed, this field will be 00h.

Table 127: Log Sense Page 10, Extended Segment Number

Extended Segment Number	Short Self-Test	Extended Self-Test
1h	Drive Ready Test	
2h	Drive Diagnostics	
3h	SMART	
4h	Low Level Format check	
5h	PLI Capacitor Self-Test	
6h	Random Verify	
7h	- Verify First 300 MB - Verify Last 100 MB	Verify all LBAs
8h	Recheck SMART	

- **Timestamp** This field contains the total accumulated power-on hours of the Target at the time the self-test completed.
- **LBA of first failure** This field contains the LBA of the first logical block address where a self-test error occurred. If no errors occurred during the self-test or the error is not related to a LBA then the field will be FFFFFFFFFFFFFFFFh.
- **Sense Key, Additional Sense Code and Additional Sense Code Qualifier** These fields will contain the additional information relating to the error or exception conditions during self-test.

See Section 17.40 “SEND DIAGNOSTIC (1D)” on page 232, for detailed listing of operations carried out by SEND DIAGNOSTIC command and Power on Diagnostics.

17.7.11 Solid State Media log page (Page 11)

The Solid State Media log page indicates parameters that are specific to SCSI target devices that contain solid state media. A device server that implements the Solid State Media log page shall implement one or more of the defined parameters.

Table 128: Solid State Media log page

Byte	Bit							
	7	6	5	4	3	2	1	0
0	DS	SPF (0)	Page code = 11h					
1	Reserved							
2-3	Page Length = (n -3)							
	Solid State Media log parameters							
4	Solid State Media parameter (first) see Table 129:							
...	First Medium Scan Parameter							
...	...							
n	Solid State Media parameter (last) see Table 129:							

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.

Table 129: Solid State Media log page parameter codes

Parameter Code	Description
0001h	Percentage Used Endurance Indicator
All others values	Reserved

Table 130: Percentage Used Endurance Indicator parameter format

Byte	Bit							
	7	6	5	4	3	2	1	0
0-1	PARAMETER CODE (0001h)							
2	DU	Obsolete	TSD	ETC	TMC	Formant and Linking (11b)		
3	PARAMETER LENGTH (04h)							
4 -6	Reserved							
7	PERCENTAGE USED ENDURANCE INDICATOR							

The FORMAT AND LINKING field shall be set to 11b, indicating that this parameter is a binary format list parameter. The values for the other bits and fields in the parameter control byte for a binary format list parameter are defined in SPC-4.

The PARAMETER LENGTH field indicates the number of bytes to follow in the log parameter.

The PERCENTAGE USED ENDURANCE INDICATOR field indicates an estimate of the percentage of device life 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 device has been consumed, but may not indicate a device failure (e.g., minimum power-off data retention capability reached for devices 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

17.7.12 Log Sense Page 15

This page contains information about Background Medium Scan operations.

Table 131: Log Sense Page 15

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Reserved		Page code = 15h					
1	Reserved							
2-3	Page Length = (19 + 24N -3)							
	Background Medium Scan parameters							
4-19	BMS Status Parameter							
20-43	First Medium Scan Parameter							
	...							
19+24N	Last Medium Scan Parameter							

The following table describes the BMS Status Parameter structure.

Byte	Bit							
	7	6	5	4	3	2	1	0
0-1	Parameter Code = 0000h							
2	DU=0	DS=0	TSD=0	ETC=0	TMC=0		F&L = 11b	
3	Page Length = 0Ch							
4-7	Power On Minutes							
8	Reserved = 0							
9	BMS Status							
10-11	Number of Scans Performed							
12-13	Medium Scan Progress							
14-15	Reserved = 0							

- **Power On Minutes** indicates the total power on minutes at the time the log page is requested
- **BMS Status** is described in the following table

BMS Status	Description
00h	No scans active
01h	Background medium scan is active
02h	Background pre-scan is active

BMS Status	Description
03h-04h	Not supported
05h	Background scan halted due to medium formatted without P-List
06h	Background scan halted due to a vendor-specific cause
07h	Background scan halted due to temperature out of range
08h	Scan suspended until BMS Interval Timer expires
09h - FFh	Reserved

- Number of Scans Performed indicates the number of background scans that have been performed over the life of the drive.
- Medium Scan Progress is a percent complete indication of the medium scan. The returned value is a numerator that has 65,536 (1 00 00h) as its denominator.

The following table describes the Medium Scan Parameter structure.

Byte	Bit							
	7	6	5	4	3	2	1	0
0-1	Parameter Code = 0001h - 0800h							
2	DU=0	DS=0	TSD=0	ETC=0	TMC=0		F&L = 11b	
3	Page Length = 14h							
4-7	Power On Minutes							
8	Reassign Status				Sense Key			
9	Additional Sense Code							
10	Additional Sense Code Qualifier							
11-15	Reserved = 0							
16-23	LBA							

- **Power On Minutes** indicates the total power on minutes at the time the error was detected.
- **Reassign Status** is set to 0h. Auto-reallocation is automatic and no action needs to be taken by the Initiator

17.7.13 Log Sense Page 1A

Table 132: Log Sense Page 1A

Byte	Bit								
	7	6	5	4	3	2	1	0	
0	Reserved		Page code = 1Ah						
1	SubPage Code (00h)								
2-3	Page Length (30h)								
4-5	Parameter Code 0001h								
6	DU	Obsolete	TSD	ETC	TMC	TMC	FMT & Linking	FMT & Linking	
7	Parameter Length = 4								
8-11	Accumulated Transitions to Active State								
12-13	Parameter Code 0002h								
14	DU	Obsolete	TSD	ETC	TMC	TMC	FMT & Linking	FMT & Linking	
15	Parameter Length = 4								
16-19	Accumulated Transitions to Idle_A								
20-21	Parameter Code 0003h								
22	DU	Obsolete	TSD	ETC	TMC	TMC	FMT & Linking	FMT & Linking	
23	Parameter Length = 4								
24-27	Accumulated Transitions to Idle_B								
28-29	Parameter Code 0004h								
30	DU	Obsolete	TSD	ETC	TMC	TMC	FMT & Linking	FMT & Linking	
31	Parameter Length = 4								
32-35	Accumulated Transitions to Idle_C								
36-37	Parameter Code 0008h								
38	DU	Obsolete	TSD	ETC	TMC	TMC	FMT & Linking	FMT & Linking	
39	Parameter Length = 4								
40-43	Accumulated Transitions to Standby_Z								
44-45	Parameter Code 0009h								
46	DU	Obsolete	TSD	ETC	TMC	TMC	FMT & Linking	FMT & Linking	
47	Parameter Length = 4								
48-51	Accumulated Transitions to Standby_Y								

17.7.14 Log Sense Page 2F

This page contains SMART Status and Temperature Reading.

Table 133: Log Sense Page 2F

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Reserved		Page code = 2Fh					
1	Reserved							
2-3	PageLength = 2Ch							
4-5	Parameter Code = 0000h							
6	DU = 0	DS = 0	TSD = 0	ETC = 0	TMC = 0		F&L = 11b	
7	Parameter Length = 04h							
8	SMART Sense Code Byte							
9	SMART Sense Qualifier							
10	Most Recent Temperature Reading							
11	Vendor Temperature Trip Point							
12	Vendor Unique Maximum Temperature							
13-15	Vendor Unique Reserved = 000000h							
16-47	Vendor Unique Parameters (see Table 134:)							

Table 134: Vendor Unique parameter Code = 0000

Byte	Bit							
	7	6	5	4	3	2	1	0
0-1	Parameter code =0000h							
1	DU = 0	DS = 0	TSD = 0	ETC = 0	TMC = 0		F&L = 11b	
2-3	PageLength = 04h							
4-5	SMART Parameter Sense Code Byte							
6	SMART Parameter Sense Qualifier							
7	SMART Attribute Percentage of Threshold							
8	SMART Attribute Trip							

SMART PARAMETER SENSE CODE is the one-byte value indicating the severity of this particular parameter when host notification for SMART trip is made. For example, 0x5D indicates pre-fail attribute and 0x0B indicates warning attribute.

SMART PARAMETER SENSE QUALIFIER is the one-byte value that uniquely identifies each particular parameter when host notification for SMART trip is made.

SMART ATTRIBUTE PERCENTAGE OF THRESHOLD indicates an estimate of the percentage of threshold reached for the vendor unique SMART attributes. The value in the field is set to zero at the time of manufacture. A value of 100 indicates that the threshold has been reached and SMART trip will be reported to the host if enabled. See Mode Page 0x1C (Information Exceptions Control). The value is allowed to exceed 100. Values greater than 254 are reported as 255. The device server shall update the value at least once per power-on hour. Note that the Volatile memory backup attribute is a pass/fail indicator so it will always read 0 unless the capacitor self test fails, and in that case it would report 100.

SMART ATTRIBUTE TRIP is set to 1b if the threshold for that SMART attribute has ever been exceeded. It is set to 0b if the threshold has never been exceeded.

For the vendor unique parameters codes > 0000h:

- Parameter code = 0001h Remaining Reserve 1
- Parameter code = 0002h Remaining Reserve 2
- Parameter code = 0003h Volatile Memory Backup Failure
- Parameter code = 0004h Wear Indicator
- Parameter code > 0004h Reserved

17.7.15 Log Sense Page 30

Table 135: Log Sense Page 30

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Reserved		Page code = 30h					
1	Reserved = 0							
2-3	Page Length = 0030h							
4-5	Parameter Code = 0000h							
6	DU = 0	DS = 0	TSD = 0	ETC = 0	TMC = 0		F&L = 00b	
7	Parameter Length = 2Ch							
8-51	Reserved = 0							

17.7.16 Log Sense Page 37

This page contains a series of miscellaneous data counters including information about predictive failure analysis occurrences.

Table 136: Log Sense Page 37

Byte	Bit								
	7	6	5	4	3	2	1	0	
0	Reserved		Page code = 37h						
1	Reserved								
2-3	Page Length = 0030h (48)								
4-5	Parameter Code = 0000h								
6	DU=0	DS=0	TSD=0	ETC=0	TMC = 0		F&L = 00b		
7	Parameter Length = 2Ch								
8	(MSB)								
-	Power on Hours (hours only)								
11	(LSB)								
12	(MSB)								
-	Total Bytes Read								
19	(LSB)								
20	(MSB)								
-	Total Bytes Written								
27	(LSB)								
28	Max Drive Temp (degrees Celsius)								
29 - 30	(MSB) Reserved = 0								
	(LSB)								
31	Number of Information Exceptions								
32	MED EXC	HDW EXC	Reserved = 0						
33 - 40	Total Read Commands								
41 - 48	Total Write Commands								
49	Reserved = 0								
50-51	Flash Correction Count								

The **Power on Hours** field specifies the total time the drive has been powered on in hours only.

The **Max. Drive Temperature** field specifies the maximum temperature, in degrees Celsius, the drive has ever reached.

The **Number of Information Exceptions** field gives the number of Information Exceptions during the life of the drive and not the number of Information Exceptions that have been reported. The number of reported Information Exceptions may be less due to the settings of Mode Page 0x1C. NOTE: This field does not include occurrences of any Information Exception Warnings.

If set, the **Media Exception and Hardware Exception** bits indicate that an Information Exception has occurred during the life of the drive. These flags are set during an Information Exception that may or may not coincide with the reporting of an Information Exceptions as mentioned above.

Total Read Commands counter is incremented for each Read (6) and Read (10) command received.

Total Write Commands counter is incremented for each Write (6), Write (10), Write Verify and Write Verify (16) command received.

Flash Correction Count is incremented each time ECC correction is applied to data stored in Flash ROM.

17.8 MODE SELECT (15)

Table 137: Mode Select (15)

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Command Code = 15h							
1	Reserved = 0			PF=1	Reserved = 0			SP
2	Reserved = 0							
3								
4	Parameter List Length							
5	VU = 0		Reserved = 0			FLAG	LINK	

The MODE SELECT (15) command provides a means for the Initiator to specify LUN or device parameters to the Target. It also allows an Initiator to specify options the Target uses in error recovery, caching, and formatting.

There is a single set of Mode Page parameters shared by all Initiators.

- **PF** A PF (Page Format) bit value of one indicates that the data sent by the Initiator after the Mode Select Header and the Block Descriptor, if any, complies to the Page Format. The Target ignores this field since it only accepts mode parameters in the Page Format.
- **SP** Save Pages. This indicates
 - 0 The drive shall not save the pages sent during the Data Out phase but will use them for all following commands until the power is removed, a reset is received, or a new MODE SELECT command is received.
 - 1 The drive will save the data in the reserved area of the media. It will be used for all the following commands until another MODE SELECT command is issued. This information is maintained over a power cycle or reset of the drive.
- **Parameter List Length** This specifies the number of bytes to be sent from the Initiator. A parameter list length of zero suppresses data transfer and is not considered an error.

The MODE SELECT parameter list contains a 4-byte header followed by zero or one block descriptor followed by zero or more pages. The pages that are valid with this command are defined in the addendum under the heading **Mode Select Data**, as they vary with the drive model.

Application Note

The Initiator should issue a MODE SENSE command requesting all Changeable values (see PCF field in byte two of the CDB) prior to issuing a MODE SELECT command. This is necessary to find out which pages are implemented by the drive and the length of those pages. In the Pages of the MODE SENSE command the drive will return the number of bytes supported for each Page. The Page Length set by the Initiator in the MODE SELECT command must be the same value as returned by the drive in MODE SENSE Page Length. If not, the drive will return *Check Condition* status with sense key of *Illegal Request*.

Note: If an Initiator sends a MODE SELECT command that changes any parameters that apply to other Initiators, the drive shall generate an unit attention condition for all Initiators except for the one that issued the MODE SELECT command. The drive shall set the additional sense code to *Parameters Changed* (2Ah).

17.9 MODE SELECT (55)

Table 138: Mode Select (55)

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Command Code = 55h							
1	Reserved = 0			PF=1	Reserved = 0			SP
2-6	Reserved = 0							
7-8	(MSB) Parameter List Length (LSB)							
9	VU = 0		Reserved = 0			FLAG	LINK	

The MODE SELECT (55) command provides a means for the Initiator to specify LUN or device parameters to the Target. See the MODE SELECT (15) command for a description of the fields in this command.

17.10 MODE SENSE (1A)

Table 139: Mode Sense (1A)

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Command Code = 1Ah							
1	Reserved			RSVD	DBD	Reserved = 0		
2	PCF		Page Code					
3	Subpage Code							
4	Allocation Length							
5	VU = 0		Reserved = 0			FLAG	LINK	

The MODE SENSE (1A) command provides a means for the drive to report various device parameters to the Initiator. It is the complement to the MODE SELECT command.

If the **DBD** (Disable Block Descriptor) bit is zero, the Target will return the Block Descriptor. If the DBD bit is set to 1, the Target will not return the Block Descriptor.

Allocation Length indicates the maximum number of bytes that the Initiator has set aside for the DATA IN phase. A value of zero is not considered an error. If the allocation length is smaller than the amount available, that portion of the data up to the allocation length will be sent. This may result in only a portion of a multi-byte field being sent.

Page Control Field: PCF (Page Control Field) defines the type of Page Parameter values to be returned.

PCF Meaning

0 0 Report current values. The drive returns the current values under which the logical unit is presently configured for the page code specified. The current values returned are

1. Initially following power-up but before the media is accessed, the default values become current. Once the media can be accessed, the saved values are read from the Reserved Area and become current.
2. The parameters set in the last successful MODE SELECT command.
3. The saved values if a MODE SELECT command has not been executed since the last power-on, hard RESET condition, or TARGET RESET message.

Following the completion of start-up, execution of the MODE SELECT command can modify the current values.

Note: Those parameters associated with format are not considered current and are not saved until the successful completion of a FORMAT UNIT command.

In addition, the current values take on the saved values after a reset if the parameters were saved. If the Page Code is 3Fh, then all pages implemented by the Target are returned to the Initiator with fields and bit values set to current values.

If the Page Code is not 3Fh, the page defined by the Page Code, if supported by the Target, is returned with fields and bits set to current values.

Note: The drive will not process the MODE SELECT command until the completion of spin-up. Therefore, the Initiator cannot modify the current values prior to the saved values being read in.

0 1 Report changeable value. The drive returns the changeable values for the page code specified. The page requested is returned containing information that indicates which fields are changeable. All bits of parameters that are changeable shall be set to one. Parameters that are *defined by the drive* shall be set to zero. If any part of a field is changeable, all bits in that field shall be set to one.

Note: For a value field such as the buffer ratios of page 2 the bit field will not indicate the range of supported values but rather that the field is supported.

1 0 Report default value. The drive returns the default values for the page code specified. The parameters not supported by the drive are set to zero.

1 1 Report saved value. The drive returns the saved value for the page code specified.

Saved values are one of the following:

- the values saved as a result of MODE SELECT command
- identical to the default values
- zero when the parameters are not supported

The Page Length byte value of each page returned by the drive indicates up to which fields are supported on that page.

Page Code: This field specifies which page or pages to return. Page code usage is defined in the figure below.

Table 140: Page Code Usage

Page Code	Description
00h - 1Ch	Return specific page, if supported.
3Fh	Return all supported pages.

If a Page Code of 3Fh is used, MODE SENSE returns the pages in ascending order with one exception. Page 0 is always returned last in response to a MODE SENSE command.

If an unsupported page is selected, the command is terminated with a CHECK CONDITION status and available sense of ILLEGAL REQUEST/INVALID FIELD IN CDB.

Subpage Code: This field specifies the subpage to return, and may be set to a specific page, or to FFh for all supported subpages.

17.10.1 Mode Parameter List

The mode parameter list contains a header followed by zero or more block descriptors followed by zero or more variable length pages.

17.10.1.1 Header

The header used for the 6-byte CDB is defined below

Table 141: Mode parameter header (6)

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Mode Data Length							
1	Medium Type = 0							
2	WP=0	Reserved=0		DPOFUA=1	Reserved = 0			
3	Block Descriptor Length (= 0 or 8)							

The header used for the 10-byte CDB is defined below.

Table 142: Mode parameter header (10)

Byte	Bit							
	7	6	5	4	3	2	1	0
0 1	Mode Data Length (MSB) (LSB)							
2	Medium Type = 0							
3	WP=0	Reserved=0	DPOFUA =1	Reserved = 0				
4 5	Reserved = 0							
6 7	Block Descriptor Length (= 0 or 8) (MSB) (LSB)							

- **Mode Data Length.** When using the MODE SENSE command, the mode data length field specifies the length in bytes of the following data that is available to be transferred. The mode data length does not include the length byte itself. When using the MODE SELECT command, this field is reserved.
- **Medium Type** field is always set to zero in the drive (Default Medium Type).
- **WP.** When used with the MODE SELECT command, the Write Protect (WP) bit is reserved. When used with the MODE SENSE command, a Write Protect (WP) bit of zero indicates that the medium is write enabled.
- **DPOFUA** bit value of 1 indicates that the Target supports the FUA and DPO bits in the Read and Write Commands.
- **Block Descriptor Length** specifies the length in bytes of the block descriptors. When used with the MODE SELECT command, zero or eight is supported by the drive. When used with the MODE SENSE command, the drive returns eight to indicate that only a single block descriptor is available.

Note: DPOFUA is ignored during Mode Select command processing although the SCSI Standard states that it is reserved during Mode Select. Ignoring it allows the Mode Sense Parameter List for the byte containing this bit to be re-used as a Mode Select Parameter List.

17.10.1.2 Block Descriptor

Table 143: Mode Parameter Block Descriptor

Byte 0	(MSB)	Number of Blocks
Byte 1		
Byte 2		
Byte 3	(LSB)	
Byte 4		Density code = 0
Byte 5	(MSB)	Block Length
Byte 6		
Byte 7	(LSB)	

The Block descriptor provides formatting information about the Number of Blocks (user addressable) to format at the specified Block Length.

- Number of Blocks

When used with the MODE SELECT command, the **Number of Blocks** field must be

- Zero to indicate not to change available blocks
- 0xFFFFFFFF to indicate all available blocks
- The exact number of blocks in the data area of the drive, which can be obtained with the MODE SENSE
- The number of blocks less than exact one, in order to **CLIP** the number of blocks

Any other value is invalid and causes the command to fail with *Check Condition* status.

When used with the MODE SENSE command, the field contains the exact number of blocks.

- Density Code

- Always 0 for direct access devices.

- Block Length

The Block Length field reflects the number of bytes of user data per sector (not including any protection information). When used with the MODE SELECT command, the **Block length** field must contain the value from 512 to 528 (8 bytes step) or zero. Otherwise the drive will terminate the command with *Check Condition* status.

A FORMAT UNIT command is required to cause these parameters to become current only if the block length parameter is different from the current block length.

When used with the MODE SENSE command, the field is dependent on how the media is currently formatted.

17.10.1.3 Page Descriptor

Table 144: Mode Parameter Page Format

Byte 0	PS	SPF	Page Code
Byte 1	Page Length		
Byte 2-n	Mode Parameters		

Each mode page contains a page code, a page length, and a set of mode parameters.

When using the MODE SENSE command, a Parameter Savable (PS) bit of one indicates that the mode page can be saved by the drive in the reserved area of the drive. A PS bit of zero indicates that the supported parameters cannot be saved. When using the MODE SELECT command, the PS bit is reserved (zero).

SPF (Sub-Page Format) is set to zero to indicate the short page format is used. The bit is set to one to indicate the long format is used, supporting sub-pages. The drive supports the following mode page codes:

Table 145: Mode Parameter Page Format

Page	Description	PS
00	Vendor Unique Parameters	1
01	Read-Write Error Recovery Parameters	1
02	Disconnect/Reconnect Control Parameters	1
03	Format Device Parameters	0
04	Rigid Disk Geometry Parameters	0
07	Verify Error Recovery Parameters	1
08	Caching Parameters	1
0A	Control Mode Page	1
0C	Notch Parameters	1
19	Port Control Page	1
1A	Power Control Parameters	1
1C	Informational Exceptions Control	1

The page length field specifies the length in bytes of the mode parameters that follow. If the Initiator does not set this value to the value that is returned for the page by the MODE SENSE command, the drive will terminate the command with *Check Condition* status.

17.10.2 Mode Page 00 (Vendor Unique Parameters)

Table 146: Vendor Unique Parameters - Page 00

Byte	Bit								Default	
	7	6	5	4	3	2	1	0		
0	PS	0	Page Code = 00h							80h
1	Page Length = 0Eh								0Eh	
2	Ignored								00h	
3	Ignored								00h	
4	Ignored								00h	
5	Ignored		FDD	Ignored		CAEN	Ignored		02h	
6	Ignored			OCT (high nibble)					00h	
7	Overall Command Timer (low byte)								00h	
8	Ignored								00h	
9	Temperature Threshold								00h	
10	Command Aging Limit (Hi byte)								00h	
11	Command Aging Limit (Low byte)								30h	
12	Error Injection	Read Reporting Threshold							16h	
13-14	Ignored								00h	
15	Ignored		FCERT	Ignored		SRS	Reserved = 0		00h	

Fields marked in the table as 'Ignored' are not used or checked by the drive. They will be initialized to zero but can be set as desired for compatibility with older drives.

- **FDD** (Format Degraded Disable) controls the reporting of Format Degraded sense data for Test Unit Ready commands when the drive is in a format degraded state. When the FDD bit is one, Format Degraded sense data will not be reported for a Test Unit Ready command. When the FDD bit is zero, Format Degraded sense data will be reported for Test Unit Ready commands when the drive is in a format degraded state. This bit does not affect the reporting of Format Degraded conditions for any media access commands.
- **CAEN** (Command Aging Enable) When set this bit causes the Command Age Limit timer to be used to avoid commands waiting in the command queue for an indefinite period. When commands have been in the queue for a period of time greater than the timer limit they will be reordered to be executed on a first come first served basis. When this bit is reset, commands are always executed based on the queue reordering rules.
- **OCT** (Overall Command Timer) controls the maximum command execution time, from receipt by the drive until status is returned. If the command is unable to complete in the specified amount of time, it will be aborted with Check Condition status, Aborted Command sense key. The Overall Command Timer does not alter the behavior of the Command Aging Limit or Recovery Time Limit. Each unit of this timer is 50 milliseconds. Setting the value to zero disabled the feature.
- **Temperature Threshold** Temperature Threshold specifies the threshold value in degrees Celsius for the thermal sensor Information Exception Warning; the reporting of which is controlled by Mode Page 0x1C. A value of 0 selects the default value (70 degrees Celsius).

- **Command Aging Limit** This value controls the maximum time a command should wait in the command queue when the CAEN bit is set. Each unit of this timer is 50 ms.
- **Error Injection** bit indicates whether internal error injection is currently enabled or disabled. This bit is read only. A value of 1 indicates that error injection mode is currently enabled. A value of 0 indicates that error injection mode is currently disabled. Error injection mode is a drive feature that randomly injects pseudo errors during read commands. It is controlled by the DPRY bit of the Format command. This mode should only be used in a drive test mode to validate drive and system error reporting and handling functionality.
- **Read Reporting Threshold** specifies the bits-in-error threshold at which recovered errors will be reported when PER in Mode Page 0x01 is set to 1. The bit error on the media needs to exceed this threshold before it will be reporting. Valid values for this field range from 3 to 23 bits in error.
- **FCERT** (Format Certification) bit determines whether the certification step will be performed during a Format Unit command. FCERT bit set to 0 disables certification. FCERT bit set to 1 enables the certification step.

17.10.3 Mode Page 01 (Read/Write Error Recovery Parameters)

Table 147: Mode Page 01 (Vendor Unique Parameters)

Byte	Bit								Default	
	7	6	5	4	3	2	1	0		
0	PS	0	Page Code = 01h							81h
1	Page Length = 0Ah								0Ah	
2	AWRE	ARRE	TB	RC	EER=0	PER	DTE	DCR	C0h	
3	Read Retry Count								01h	
4	Obsolete = 0								00h	
5	Obsolete = 0								00h	
6	Obsolete = 0								00h	
7	Reserved								00h	
8	Write Retry Count								00h	
9	Reserved								00h	
10	(MSB) Recovery Time Limit (LSB)								00h	
11										

The Read-Write recovery parameters that will be used during any command that performs a read or write operation to the medium are as follows:

- **AWRE** Automatic Write Reallocation Enabled bit, is ignored. Automatic Write Reallocation is always performed
- **ARRE** Automatic Read Reallocation Enabled bit is ignored. Automatic Read Reallocation is always performed.
- **TB** (Transfer Block bit) is ignored.
- **RC** (Read Continuous) bit, is ignored.
- **EER** (Enable Early Recovery) bit is ignored.
- **PER** (Post Error) bit specifies whether or not recovered errors are reported. It is used in conjunction with the Read Reporting Threshold in Mode Page 0x00. Setting this bit to 1 enables recovered error reporting of media bit errors that exceed the Read Reporting Threshold in Page 0x00. Setting this bit to 0 disables recovered error reporting of media errors.
- **DTE** (Data Terminate on Error) is ignored.
- **DCR** (Disable Correction) bit, is ignored.
- **Read Retry Count** ignored. Read recovery is always performed.
- **Write Retry Count** s ignored.
- **Recovery Time Limit** is ignored

17.10.4 Mode Page 02 (Disconnect/Reconnect Parameters)

Table 148: Mode Page 02 (Disconnect/Reconnect Parameters)

Byte	Bit								Default
	7	6	5	4	3	2	1	0	
0	PS	0	Page Code = 02h						82h
1	Page Length = 0Eh								0Eh
2	Read Buffer Full Ratio								00h
3	Write Buffer Empty Ratio								00h
4-5	(MSB)	Bus Inactivity Limit = 0						(LSB)	00h
6-7	(MSB)	Disconnect Time Limit = 0						(LSB)	00h
8-9	(MSB)	Maximum Connect Time Limit						(LSB)	00h
10-11	(MSB)	Maximum Burst Size						(LSB)	00h
12-15	Reserved = 0								00h

The disconnect/reconnect page provides the Initiator with the means to tune the performance of the Fibre Channel Loop.

The drive uses the disconnect/reconnect parameters to control when it attempts to regain control of the Loop during READ and WRITE commands.

- **Read Buffer Full Ratio** is ignored.
- **Write Buffer Empty Ratio** is the numerator of a fraction whose denominator is 256. The fraction indicates how empty the drive data buffer should be before attempting to re-arbitrate for the Loop. If the ratio is set to 0h, the target will calculate and use an optimal ratio based on the negotiated transfer rate.

Both the Read Buffer Full Ratio and the Write Buffer Empty Ratio pertain to the current active notch. For each active notch as defined in page 0Ch there are separate Read Buffer Full Ratios and Write Buffer Empty Ratios. When the active notch is set to zero, the values are applied in mode page 0Ch across all notches.

- **Maximum Connect Time Limit** specifies the maximum amount of time the drive will keep a FCAL connection open. The time is specified in 100 microsecond increments. The default value of zero, indicates no time limit. A maximum value of FFFFh, specifies a connection time limit of 6.55 seconds. When this time expires, the drive will prepare to close the connection.
- **Maximum Burst Size field** indicates the maximum amount of data that the target port shall transfer during a single data transfer operation. This value is expressed in increments of 512 bytes. A value of zero specifies there is no limit on the amount of data transferred per data transfer operation.

17.10.5 Mode Page 03 (Format Device Parameters)

Table 149: Mode Page 03 (Format Device Parameters)

Byte	Bit								Default	
	7	6	5	4	3	2	1	0		
0	PS	0	Page Code = 03h							03h
1	Page Length = 16h								16h	
2-3	(MSB) Tracks per Zone								00h	
	(LSB)								00h	
4-5	(MSB) Alternate Sectors per Zone = 0								00h	
	(LSB)								00h	
6-7	(MSB) Alternate Tracks per Zone = 0								00h	
	(LSB)								00h	
8-9	(MSB) Alternate Tracks per Logical Unit = 0								00h	
	(LSB)								00h	
10-11	(MSB) Sectors Per Track								00h	
	(LSB)								00h	
12-13	(MSB) Data Bytes per Physical Sector								00h	
	(LSB)								00h	
14-15	(MSB) Interleave = 0001h or 0000h								00h	
	(LSB)								01h	
16-17	(MSB) Track Skew Factor								00h	
	(LSB)								00h	
18-19	(MSB) Cylinder Skew Factor								00h	
	(LSB)								00h	
20	SSEC	HSEC	RMB	SURF	RESERVED = 0				40h	
21-23	Reserved = 0								00h	

The format device page contains parameters that specify the medium format. This page contains no changeable parameters.

- **Tracks per Zone** is obsolete for SSDs.
- **Sectors per Track** is obsolete for SSDs.
- **Data Bytes per Physical Sector** specifies the number of user data bytes per physical sector. The value depends upon the current formatted Block Length.
- **Interleave** value of 1 or 0 is valid. However, the drive will ignore this.
- **Track Skew Factor** is obsolete for SSDs.
- **Cylinder Skew Factor** is obsolete for SSDs.
- **SSEC = Zero** indicates that the drive does not support soft sector formatting.
- **HSEC = One** indicates that the drive supports hard sector formatting.
- **RMB = Zero** indicates that the media does not support removable Fixed drive.
- **SURF = Zero** indicates that progressive addresses are assigned to all logical blocks in a cylinder prior to allocating addresses within the next cylinder.

17.10.6 Mode Page 04 (Rigid Disk Drive Geometry Parameters)

Table 150: Mode Page 04 (Rigid Disk Drive Geometry Parameters)

Byte	Bit								Default
	7	6	5	4	3	2	1	0	
0	PS	0	Page Code = 04h						04h
1	Page Length = 16h								16h
2-4	(MSB) Number of Cylinders								00h
	(LSB)								00h
5	Number of heads								00h
6-8	(MSB) Starting Cylinder - Write Precompensation = 0								00h
	(LSB)								00h
9-11	(MSB) Starting Cylinder - Reduced Write Current = 0								00h
	(LSB)								00h
12-13	(MSB) Drive Step Rate = 0 (Not used)								00h
	(LSB)								00h
14-16	(MSB) Landing Zone Cylinder = 0 (Not used)								00h
	(LSB)								00h
17	RESERVED = 0					RPL = 0			00h
18	Rotational Offset = 0 (Not used)								00h
19	RESERVED = 0								00h
20-21	(MSB) Medium Rotation Rate								00h
	(LSB)								01h
22-23	Reserved = 0								00h

The rigid disk drive geometric page specifies various parameters for the drive.

- Medium Rotation Rate = 1 indicates the drive is an SSD.

17.10.7 Mode Page 07 (Verify Error Recovery Parameters)

Table 151: Mode Page 07 (Verify Error Recovery Parameters)

Byte	Bit								Default
	7	6	5	4	3	2	1	0	
0	PS	0	Page Code = 07h						87h
1	Page Length = 0Ah								0Ah
2	Reserved = 0				EER=0	PER	DTE	DCR	00h
3	Verify Retry Count								01h
4	Obsolete =0								00h
5 - 9	Reserved = 0								00h
10-11	(MSB)	Verify Recovery Time Limit						(LSB)	00h

The Verify recovery parameters are used by the Target when recovering from and reporting errors associated with the verification of the Initiator's Data for the following commands:

- **VERIFY**
- **WRITE AND VERIFY** - the verify portion of the command only.
- **EER.** is ignored.
- **PER.** is ignored. The PER setting in Mode Page 01h is used for Verify commands.
- **DTE.** is ignored.
- **DCR.** is ignored.
- **Verify Recovery Time Limit** is ignored.

17.10.8 Mode Page 08 (Caching Parameters)

Table 152: Page 08 (Caching Parameters)

Byte	Bit								Default
	7	6	5	4	3	2	1	0	
0	PS	0	Page Code = 08h						88h
1	Page Length = 12h								12h
2	IC	ABPF	CAP	DISC	SIZE	WCE	MF	RCD	04h
3	Demand Read Retention Priority				Write Retention Priority				00h
4-5	(MSB) Disable Pre-fetch Transfer Length (LSB)								FFh
6-7	(MSB) Minimum Pre-fetch (LSB)								00h
8-9	(MSB) Maximum Pre-fetch (LSB)								FFh
10-11	Maximum Pre-fetch Ceiling								FFh
12	FSW	LBCSS	DRA	Reserved = 0					00h
13	Number of Cache Segments								08h
14-15	(MSB) Cache Segment Size (LSB)								00h
16	Reserved = 0								00h
17-19	(MSB) Non Cache Segment Size (LSB)								00h

The caching parameters page defines parameters that affect the use of the cache.

- **IC** (Initiator Control) bit is ignored.
- **ABPF** (Abort Pre-fetch) bit is ignored.
- **CAP** (Caching Analysis Permitted) is not supported and is ignored.
- **DISC** (Discontinuity) is not supported and is ignored.
- **SIZE** (Size Enable) bit is ignored.
- **WCE** (Write Cache Enable) bit is ignored. A fail-safe write caching is always enabled.
- **MF** (Multiplication Factor) bit is ignored.
- **RCD** (Read Cache Disable) bit is ignored.
- **Demand Read Retention Priority** is not supported.
- **Write Retention Priority** is not supported.
- **Disable Pre-fetch Transfer Length** is ignored.
- **Minimum Pre-fetch** is ignored.
- **Maximum Pre-fetch** is ignored.
- **Maximum Pre-fetch Ceiling** is ignored.
- **FSW** (Force Sequential Write) is not supported and is ignored. All logical blocks will be written in sequential order.
- **LBCSS** (Logical Block Cache Segment Size) bit is ignored
- **DRA** (Disable Read Ahead) is ignored.

- **Number of Cache Segments** field is ignored.
- **Cache Segment Size** field is ignored.
- **Non Cache Segment Size** is not supported and is ignored

17.10.9 Mode Page 0A (Control Mode Page Parameters)

Table 153: Page 0A (Control Mode Page Parameters)

Byte	Bit								Default	
	7	6	5	4	3	2	1	0		
0	PS	0	Page Code = 0Ah							8Ah
1	Page Length = 0Ah								0Ah	
2	TST=0		TMFonly=0	RSVD=0	D_Sense=0	GLTSD=0	RLEC=0		00h	
3	Queue Algorithm Modifier			Rsvd=0	QErr		DQue		00h	
4	RSVD=0	RAC=0	UA_INTLCK_CTRL=0	SWP=0	Obsolete				00h	
5	ATO	TAS=0	Reserved=0						00h	
6-7	Obsolete=0								00h	
8-9	(MSB) Busy Timeout Period (LSB)								00h	
10-11	(MSB) Extended Self-test Routine Completion Time (LSB)								xxh	

Following are parameter options for Page 0A.

- **Queue algorithm modifier** specifies restrictions on the algorithm used for reordering commands that are tagged with the SIMPLE message.
 - 0h: Restricted reordering. The Target shall reorder the actual execution sequence of the queued commands from each Initiator such that data integrity is maintained for that Initiator.
 - 1h: Unrestricted reordering allowed. The Target may reorder the actual execution sequence of the queued commands in any manner it selects. Any data integrity exposures related to command sequence order are explicitly handled by the Initiator through the selection of appropriate commands and queue tag messages.
 - 2h-7h: RESERVED.
 - 8: Command reordering is disabled
 - 9-Fh: RESERVED
- **QErr** (Queue Error Management) The queue error management (QERR) field specifies how the device server shall handle blocked tasks when another task receives a *Check Condition* status.

QERR value	Description
00b	Specifies that all tasks from all Initiators are blocked from execution when a Contingent Allegiance (CA condition) is pending. Those blocked tasks are allowed to resume execution in a normal fashion after the CA condition is cleared.
01b	Specifies that all tasks from all Initiators are aborted when the Target returns <i>Check Condition</i> status. A unit attention condition will be generated for each Initiator that had commands in the queue except for the Initiator that received the <i>Check Condition</i> status. The sense key will be set to <i>Unit Attention</i> and the additional sense code will be set to <i>Commands Cleared by Another Initiator</i> .

10b	Reserved
11b	Blocked tasks in the task set belonging to the Initiator to which a <i>Check Condition</i> status is sent shall be aborted when the status is sent.

- **DQue** (Disable Queuing) bit set at zero specifies that tagged queuing shall be enabled if the Target supports tagged queuing. A DQue bit set at one specifies that tagged queuing shall be disabled. Command queuing is always enabled on the drive, therefore this bit is ignored.
- **ATO** (Application Tag Owner) bit set to one specifies that the contents of the Logical Block Application Tag field in the protection information, if any, shall not be modified by the drive. An ATO bit set to zero specifies that the contents of the Logical Block Application Tag field in the protection information, if any, may be modified by the drive. If the ATO bit is set to zero, the drive will ignore the contents of the Logical Block Application Tag field in the protection information.
- **Busy Timeout Period** is not supported and is ignored.
- **Extended Self-test Routine Completion Time** is an advisory parameter that an Initiator may use to determine the time in seconds that the Target requires to complete self-test routine when the Target is not interrupted by an Initiator and no errors occur during execution of the self-test routine.

17.10.9.1 Control Extension Subpage

Table 154: Control Extension Subpage

CByte	Bit								Default
	7	6	5	4	3	2	1	0	
0	PS	SPF=1	Page Code = 0Ah						4Ah
1	Subpage Code = 1								01h
2-3	Page Length = 001Ch								001Ch
4	Reserved = 0				TCMOS	SCSIP	IALUAE	00h	
5	Reserved = 0				Initial Priority			00h	
6-31	Reserved = 0								00h

No fields in the Control Extension subpage are currently changeable. The page is supported for compatibility only.

17.10.10 Mode Page 0C (Notch Parameters)

Table 155: Page 0C (Notch Parameters)

Byte	Bit								Default
	7	6	5	4	3	2	1	0	
0	PS	0	Page Code = 0Ch					8Ch	
1	Page Length = 16h								16h
2	ND=1	LPN=0	Reserved = 0					80h	
3	Reserved = 0								00h
4-5	(MSB) Maximum Number of Notches (LSB)								00h
									00h
6-7	(MSB) Active Notch (LSB)								00h
									00h
8-11	(MSB) Starting Boundary (LSB)								00h
									00h
12-15	(MSB) Ending Boundary (LSB)								00h
									00h
16-23	(MSB) Pages Notched (LSB)								0000h
									0000h
								0000h	
								100Ch	

- The notch page is obsolete for SSDs.

17.10.11 Mode Page 19h (Fibre Channel Port Control Page)

Table 156: Mode Page 19h

Byte	BIT								Default
	7	6	5	4	3	2	1	0	
0	PS	RSVD=0	Page Code = 19h						99h
1	Page Length = 06h								06h
2	Reserved = 0								00h
3	DTFD	PLPB	DDIS	DLM	RHA	ALWLI	DTIPE	DTOLI	00h
4	Reserved = 0								00h
5									00h
6	Reserved = 0					RR_TOV Units			00h
7	Resource Recovery Time Out Value (RR_TOV)								00h

- **DTFD** (Disable Target Fabric Discovery) bit of one indicates that a Target attached by an FC-AL loop shall not recognize the presence of a fabric loop port, FL_Port, on the loop. The Target shall perform only the private loop functions defined for Targets defined by FC-PLDA. When DTFD bit is zero, the Target attached by an FC-AL loop shall discover FL_Port if present on the loop and perform the public loop functions defined for Targets by FC-FLA.
- **PLPB** (Prevent Loop Port Bypass) bit of one specifies that the Target ignores all LPB (Loop Port Bypass) and LPE (Loop Port Enable) primitive sequences. The Target's ports always remain in participating mode. A PLPB bit of zero specifies that the Target allow LPB and LPE primitive sequences to control its port bypass circuitry.
- **DDIS** (Disable Discovery) bit of one specifies the Target does not require receipt of Address or Port Discovery in order to resume tasks following loop initialization. When DDIS is zero, the Target will only resume tasks for an Initiator on receipt of an Address or Port Discovery from that Initiator.
- **DLM** (Disable Loop Master) bit of one specifies the Target does not become loop master during loop initialization. When DLM is zero, the Target may become loop master.
- **RHA** (Require Hard Address) bit of one indicates that a Target attached to an FC-AL loop shall only attempt to obtain its hard address available in the SCA-2 SFF-8067 connector or device address jumpers during loop initialization. The Target shall not attempt to obtain an address during the LISA phase of initialization. If there is a conflict for the hard address selection during loop initialization or the Target does not have a valid hard address available, the Target shall enter the non-participating state. If the Target detects loop initialization while in the non-participating state, the Target shall again attempt to get its hard address. If the hard address has not changed from the address obtained in a previous successful loop initialization, the Target shall attempt to obtain the address in the LIFA phase if a valid Fabric login exists or LIPA phase of loop initialization. If the hard address has changed, the Target shall attempt to obtain the new address in the LIHA phase. When the RHA bit is zero, the Target follows the normal initialization procedure, including the possibility of obtaining a soft address during the loop initialization process.
- **ALWLI** (Allow Login Without Loop Initialization) bit of one specifies the Target uses its hard address to accept logins without verifying the address with loop initialization. When ALWLI is zero, the Target is required to obtain an address via the loop initialization procedure before accepting a login.
- **DTIPE** (Disable Target Initiated Port Enable) bit of one specifies the Target waits for a Loop Port Enable primitive with its own hard address before inserting itself onto the loop. When DTIPE is zero, the Target inserts itself onto the loop without waiting for a Loop Port Enable primitive.

- **DTOLI** (Disable Target Originated Loop Initialization) bit of one specifies the Target does not generate the initializing LIP following insertion into the loop. The Target will respond to an initializing LIP, if received. The Target will generate the loop failure LIP if it detects a loop failure at its input and the initializing LIP when the loop failure is corrected. When DTOLI is zero, the Target generates the initializing LIP after it enables a port into a loop.
- **RR_TOV Units** field indicates the units in which the RR_TOV is calculated, according to Table 157.

Table 157: Values for RR_TOV Units

Byte 6			Units of measure for RR_TOV
bit 2	bit 1	bit 0	
0	0	0	No timer is specified
0	1	1	0.1 seconds
1	0	1	10 seconds
All other values			Reserved

- **RR_TOV** (Resource Recovery Time Out Value) field indicates the number of time units specified by the RR_TOV UNITS field that shall be used by the timer that performs the RR_TOV timeout functions. The timer accuracy is +/- 0.1 seconds. Changes to RR_TOV will only affect the usage of the timer for non-authentication situations.

When the RR_TOV Units are set to no timer, the RR_TOV value (byte 7) shall be ignored by the Target and a default timeout value of 2.0 seconds shall be used for RR_TOV.

When the RR_TOV Units are set to 0.1 seconds the following rules apply:

- a value of 00h will result in an infinite timeout value (no timeout will occur) Note: This will also cause E_D_TOV to have an infinite timeout value.
- a value from 01h to 7Eh will result in a timeout duration that is the product of the RR_TOV value and 0.1 seconds any value from 7Fh to FFh will result in the Target terminating the command with CHECK CONDITION status. The sense key is set to *Illegal Request* and the additional sense code is set to INVALID FIELD in the PARAMETER LIST.

When the RR_TOV Units are set to 10 seconds, the following rules apply:

- a value of 00h will result in an infinite timeout value (no timeout will occur) Note: This will also cause E_D_TOV to have an infinite timeout value.
- a value of 01h will result in a timeout duration that is the product of the RR_TOV value and 10 seconds
- any value from 02h to FFh will result in the Target terminating the command with CHECK CONDITION status.
- The Sense Key is set to Illegal Request and the Additional Sense Code is set to INVALID FIELD in the PARAMETER LIST.

17.10.12 Mode Page 1A (Power Control)

Table 158: Page 1A (Power Control)

Byte	Bit								Default
	7	6	5	4	3	2	1	0	
0	PS	0	Page Code = 1Ah						9Ah
1	Page Length = 26h								26h
2	Reserved = 00h						Standby_Y	00h	
3	Reserved = 0			Idle_C	Idle_B	Idle_A	Standby_Z	00h	
4-7	Idle_A Condition Timer								00h
8-11	Standby_Z Condition Timer								00h
12-15	Idle_B Condition Timer								00h
16-19	Idle_C Condition Timer								00h
20-23	Standby_Y Condition Timer								00h
24-39	Reserved								00h

- If the **STANDBY_Y** bit is set to one, then the standby_y condition timer is enabled. If the STANDBY_Y bit is set to zero, then the device shall ignore the standby_y condition timer.
- If the **IDLE_C** bit is set to one, then the idle_c condition timer is enabled. If the IDLE_C bit is set to zero, then the device shall ignore the idle_c condition timer.
- If the **IDLE_B** bit is set to one, then the idle_b condition timer is enabled. If the IDLE_B bit is set to zero, then the device shall ignore the idle_b condition timer.
- If the **IDLE_A** bit is set to one, then the idle_a condition timer is enabled. If the IDLE_A bit is set to zero, then the device shall ignore the idle_c condition timer.
- If the **STANDBY_Z** bit is set to one, then the standby_z condition timer is enabled. If the STANDBY_Z bit is set to zero, then the device shall ignore the standby_z condition timer.
- The **IDLE_A Condition Timer** field specifies the initial value, in 100 millisecond increments, for the idle_a power condition timer. The minimum allowable inactivity time for idle_a is 1 second. Any value less than this is accepted, but will automatically default to 1 second.
- The **STANDBY_Z Condition Timer** field specifies the initial value, in 100 millisecond increments, for the standby_z power condition timer. The minimum allowable inactivity time for standby_z is 2 minutes. Any value less than this is accepted, but will automatically default to two minutes. In addition, a limit of 60 timer initiated head unloads per 24 hour period is enforced.
- The **IDLE_B Condition Timer** field specifies the initial value, in 100 millisecond increments, for the idle_b power condition timer. The minimum allowable inactivity time for idle_b is 2 minutes. Any value less than this is accepted, but will automatically default to two minutes. In addition, a limit 60 timer initiated head unloads per 24 hour period is enforced.
- The **IDLE_C Condition Timer** field specifies the initial value, in 100 millisecond increments, for the idle_c power condition timer. The minimum allowable inactivity time for idle_c is 2 minutes. Any value less than this is accepted, but will automatically default to two minutes. In addition, a limit of 60 timer initiated head unloads per 24 hour period is enforced.
- The **STANDBY_Y Condition Timer** field specifies the initial value, in 100 millisecond increments, for the standby_y power condition timer. The minimum allowable inactivity time for standby_y is 2 minutes. Any value less than this is accepted, but will automatically default to two minutes. In addition, a limit 60 timer initiated head unloads per 24 hour period is enforced.

17.10.13 Mode Page 1C (Informational Exceptions Control)

Table 159: Page 1C (Informational Exceptions Control)

Byte	BIT								Default	
	7	6	5	4	3	2	1	0		
0	PS	0	Page Code = 1Ch							9Ch
1	Page Length = 0Ah								0Ah	
2	PERF	RSVD	EBF	EWASC	DEXCPT	TEST	EBACK-ERR	LOGERR	10h	
3	Reserved = 0				Method of Reporting				03h	
4-7	(MSB) Interval Timer (LSB)								00h	
									00h	
									00h	
									00h	
8-11	(MSB) Report Count (LSB)								00h	

- **PERF** (Performance) bit is not supported and is ignored. Informational Exception operations will not cause performance delays.
- **EBF** (Enable Background Function) bit is not supported and is ignored. Background functions are always enabled.
- **EWASC** (Enable Warning ASC) bit of zero indicates that Temperature Warnings will not be reported. An EWASC bit of one allows Temperature Warnings to be reported, if the temperature inside the drive enclosure exceeds the threshold set in Mode Page 00h. The Method of Reporting field controls the reporting method. EWASC is independent of DEXCPT.
- **DEXCPT** (Disable Exception Control) bit of zero indicates information exception operations are enabled. The reporting of information exception conditions when the DEXCPT bit is set to zero is determined from the Method of Reporting field. A DEXCPT bit of one indicates the Target disabled all information exception operations.
- **TEST** bit of one instructs the drive to generate false drive notifications at the next interval time, (as determined by the INTERVAL TIMER field), if the DEXCPT is zero. The Method of Reporting and Report Count would apply. The false drive failure is reported as sense qualifier 5DFFh. The TEST bit of zero instructs the drive to stop generating any false drive notifications.
- **Enable Background Error (EBACKERR)** bit of zero disables reporting of background self-test errors and background scan errors via Information Exceptions Control. An EBACKERR bit of one enables reporting of these background errors as Information Exception Warnings. The method of reporting these errors is determined from the MRIE field.
- **LOGERR** (Log Errors) is not used and ignored internally by the Target.

- **Method of Reporting** Informational Exceptions indicates the methods used by the Target to report informational exception conditions.

Code Description

- 0h** **No reporting of informational exception condition:** This method instructs the Target to not report informational exception condition.
- 1h** **Asynchronous event reporting:** Not supported.
- 2h** **Generate unit attention:** This method instructs the Target to report informational exception conditions by returning a *Check Condition* status on any command. The sense key is set to *Unit Attention* and the additional sense code indicates the cause of the informational exception condition. The command that has the *Check Condition* is not executed before the informational exception condition is reported.
- 3h** **Conditionally generate recovered error:** This method instructs the Target to report informational exception conditions, dependent on the value of the PER bit of the error recovery parameters mode page, by returning a *Check Condition* status on any command. The sense key is set to *Recovered Error* and the additional sense code indicates the cause of the informational exception condition. The command that has the *Check Condition* completes without error before any informational exception condition is reported.
- 4h** **Unconditionally generate recovered error:** This method instructs the Target to report informational exception conditions, regardless of the value of the PER bit of the error recovery parameters mode page, by returning a *Check Condition* status on any command. The sense key is set to *Recovered Error* and the additional sense code indicates the cause of the informational exception condition. The command that has the *Check Condition* completes without error before any informational exception condition is reported.
- 5h** **Generate no sense:** This method instructs the Target to report informational exception conditions by returning a *Check Condition* status on any command. The sense key is set to *No Sense* and the additional sense code indicates the cause of the informational exception condition. The command that has the *Check Condition* completes without error before any informational exception condition is reported.
- 6h** **Only report informational exception condition on request:** This method instructs the Target to preserve the informational exception(s) information. To find out about information exception conditions the Application Client polls the Target by issuing an unsolicited *Request Sense* command. The sense key is set to *No Sense* and the additional sense code indicates the cause of the informational exception condition.
- 7h-Fh** Reserved.

- **Interval Timer** field indicates the period in 100 millisecond increments for reporting that an informational exception condition has occurred. The target shall not report informational exception conditions more frequently than the time specified by the Interval Timer field and as soon as possible after the time interval has elapsed. After the informational exception condition has been reported the interval timer is restarted. A value of zero or 0xFFFFFFFF in the Interval Timer field indicates that the target only reports the informational exception condition one time and will override the value set in the Report Count field.
- **Report Count** field indicates the number of times the Target reports an informational exception condition. The Report Count of ZERO indicates no limits on the number of times the Target reports an informational exception condition.

17.10.13.1 Background Control (Subpage 01h)

Table 160: Background Control (Subpage 01h)

Byte	BIT								Default	
	7	6	5	4	3	2	1	0		
0	PS	SPF=1	Page Code = 1Ch							DCh
1	Subpage Code = 01h								01h	
2-3	Page Length = 0Ch								000Ch	
4	Reserved = 0					S_L_Full	LOWIR	EN_BMS	00h	
5	Reserved = 0							EN_PS	00h	
6-7	Background Medium Scan Interval Time								00A8h	
8-9	Background Pre-Scan Time Limit								0000h	
10-11	Minimum Idle Time Before Background Scan								0000h	
12-13	Maximum Time To Suspend Background Scan (Ignored)								0000h	
14-15	Reserved = 0								0000h	

- **Suspend On Log Full (S_L_FULL)** bit set to zero allows background scans to continue if the results log (Log Sense Page 15h) is full. S_L_FULL bit set to one will cause background scans to suspend when the log is full.
- **Log Only When Intervention Required (LOWIR)** bit set to zero allows logging of all medium errors in the results log (Log Sense Page 15h). When the LOWIR bit is set to one, only unrecovered medium errors will be logged.
- **EN_BMS (Enable Background Medium Scan)** bit set to zero specifies that the background medium scan is disabled. EN_BMS bit set to one specifies that background medium scan operations are enabled. If a background medium scan is in progress when the EN_BMS bit is changed from one to zero, then the medium scan shall be suspended until the EN_BMS bit is set to one, at which time the medium scan shall resume from the suspended location.
- **EN_PS (Enable Pre-Scan)** bit set to zero specifies that the pre-scan is disabled. If a pre-scan operation is in progress when EN_PS is changed from a one to a zero, then pre-scan is halted. An EN_PS bit set to one specifies that a pre-scan operation is started after the next power-on cycle. Once this pre-scan has completed, another pre-scan shall not occur unless the EN_PS bit is set to zero, then set to one, and another power-on cycle occurs.
- **Background Medium Scan Interval Time** specifies the minimum time, in hours, between the start of one background medium scan operation and the start of the next background medium scan operation.
- **Background Pre-Scan Time Limit** specifies the maximum time, in hours, for a pre-scan operation to complete. If the pre-scan operation does not complete within the specified time, then it is halted. A value of zero specifies an unlimited time limit.
- **Minimum Idle Time Before Background Scan** specifies the minimum time, in milliseconds, that the drive must be idle before resuming a background media scan or pre-scan. A value of zero will be treated as the default value of 1.0 second. Any value less than 100 milliseconds will be treated as 100 milliseconds. The internal timer granularity is 50 milliseconds.
- **Maximum Time To Suspend Background Scan (Ignored).**

17.11 MODE SENSE (5A)

Table 161: Mode Sense (5A)

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Command Code = 5Ah							
1	Reserved = 0				DBD	Reserved = 0		
2	PCF		Page Code					
3	Subpage Code							
4-6	Reserved = 0							
7-8	(MSB) Allocation Length (LSB)							
9	VU = 0		Reserved = 0			FLAG	LINK	

The MODE SENSE (5A) command provides a means for the drive to report various device parameters to the Initiator. See the MODE SENSE (1A) command for a description of the fields in this command.

17.12 PERSISTENT RESERVE IN (5E)

Table 162: Persistent Reserve In (5E)

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Command Code = 5Eh							
1	Reserved = 0			Service Action				
2-6	Reserved = 0							
7-8	(MSB) Allocation Length							(LSB)
9	VU = 0		Reserved = 0			FLAG	LINK	

The PERSISTENT RESERVE IN command is used to obtain information about persistent reservations and reservation keys that are active within the controller. This command is used in conjunction with the PERSISTENT RESERVE OUT command PERSISTENT RESERVE OUT (5F).

The **Allocation Length** indicates how much space has been allocated for the returned parameter data. If the length is not sufficient to contain all parameter data, the first portion of the data will be returned. If the remainder of the data is required, the initiator should send a new PERSISTENT RESERVE IN command and an Allocation Length large enough to contain all data.

17.12.1 Service Action

The following service action codes are implemented. If a reserved service action code is specified, the drive returns a **Check Condition** status. The sense key is set to *Illegal Request* and the additional sense data is set to *Invalid Field in CDB*.

Table 163: PERSISTENT RESERVE IN, Service Action Codes

Code	Name	Descriptions
00h	Read Keys	Reads all registered Reservation Keys
01h	Read Reservations	Reads all current persistent reservations
02h	Report Capabilities	Returns capability information
03h	Read Full Status	Reads complete information about all registrations and the persistent reservation, if any
04h-1Fh	Reserved	Reserved

17.12.2 Parameter data for Read Keys

Table 164: PERSISTENT RESERVE IN, parameter data for Read Keys

Byte	Bit							
	7	6	5	4	3	2	1	0
0-3	(MSB) Generation (LSB)							
4-7	(MSB) Additional length (n-7) (LSB)							
8-15	(MSB) First reservation key (LSB)							
	:							
(n-7) - n	(MSB) Last reservation key (LSB)							

Generation is a counter that increments when PERSISTENT RESERVE OUT command with “Register” or “Preempt and Clear” completes successfully. Generation is set to 0 as part of the power on reset process and hard reset process.

The **Generation** field contains a 32-bit counter that the Target shall increment every time a PERSISTENT RESERVE OUT command requests a Register, a Clear, a Preempt, or a Preempt and Abort service action. The counter shall not be incremented by a PERSISTENT RESERVE IN command, by a PERSISTENT RESERVE OUT command that performs a Reserve or Release service action, or by a PERSISTENT RESERVE OUT command that is not performed due to an error or reservation conflict. Regardless of the APTPL value the generation value shall be set to 0 as part of the power on reset process.

The **Additional Length** field contains a count of the number of bytes in the reservation key list. If the allocation length specified by the PERSISTENT RESERVE IN command is not sufficient to contain the entire parameter list, then only the bytes from 0 to the maximum allowed allocation length shall be sent to the Initiator. The incremental remaining bytes shall be truncated, although the Additional Length field shall still contain the actual number of bytes in the reservation key list without consideration of any truncation resulting from an insufficient allocation length. This shall not be considered an error.

The **Reservation Key** list contains the 8-byte reservation keys for all Initiators that have registered through all ports with the Target.

17.12.3 Parameter Data for Read Reservations

Table 165: PERSISTENT RESERVE IN, parameter data for Read Reservations

Byte	Bit							
	7	6	5	4	3	2	1	0
0-3	(MSB) Generation (LSB)							
4-7	(MSB) Additional length (n-7) (LSB)							
8-n	(MSB) Reservation descriptors (LSB)							

The **Generation** field shall be as defined for the Persistent Reserve In Read Keys parameter data. The Additional Length field contains a count of the number of bytes to follow in the Reservation Descriptor(s).

If the **Allocation length** specified by the PERSISTENT RESERVE IN command is not sufficient to contain the entire parameter list, then only the bytes from 0 to the maximum allowed allocation length shall be sent to the Initiator. The remaining bytes shall be truncated, although the Additional Length field shall still contain the actual number of bytes of the Reservation Descriptor(s) and shall not be affected by the truncation. This shall not be considered an error.

The format of the **Reservation Descriptors** is defined in the Persistent Reserve In Reservation Descriptor table. There shall be a Reservation Descriptor for the persistent reservation, if any, present in the Target having a persistent reservation.

Table 166: PERSISTENT RESERVE IN, Read Reservation Descriptor

Byte	Bit							
	7	6	5	4	3	2	1	0
0-7	(MSB) Reservation key (LSB)							
8-11	(MSB) Scope-specific address (LSB)							
12	Reserved							
13	Scope=0				Type			
14-15	(MSB) Extent Length=0 (LSB)							

The **Scope** of each persistent reservation created by a PERSISTENT RESERVE OUT command will be returned. See the PERSISTENT RESERVE OUT command section for details.

17.13 PERSISTENT RESERVE OUT (5F)

Table 167: PERSISTENT RESERVE OUT (5F)

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Command Code = 5Fh							
1	Reserved = 0			Service Action				
2	Scope=0				Type			
3-6	Reserved = 0							
7-8	Parameter List Length = 18h							
11	VU = 0		Reserved = 0				FLAG	LINK

The PERSISTENT RESERVE OUT command is used to request service actions that reserve the drive for the exclusive or shared use of the initiator. The command uses other service actions to manage and remove such reservations. This command is used in conjunction with the PERSISTENT RESERVE IN command, and should not be used with the RESERVE and RELEASE commands.

Note: If a PERSISTENT RESERVE OUT command is received when a RESERVE is active for the drive, the command will be rejected with **Reservation Conflict** status.

Parameter List Length must be 18h. If not, Check Condition status will be returned, with sense key of Illegal Request and additional sense code of Parameter List Length Error.

17.13.1 Service Action

The following service action codes are supported.

Table 168: PERSISTENT RESERVE OUT, Service Action Code

Code	Name	Description
00h	Register	Register a reservation key
01h	Reserve	Create a persistent reservation using a reservation key
02h	Release	Release a persistent reservation
03h	Clear	Clear all reservation keys and all persistent reservations
04h	Preempt	Preempt persistent reservations from another Initiator
05h	Preempt and Abort	Preempt persistent reservations from another Initiator and clear the task set for the preempted Initiator
06h	Register and Ignore existing key	Register a reservation key
07h-1Fh	Reserved	Reserved

17.13.2 Type

The **Type** field specifies the characteristics of the persistent reservation being established for all customer data sectors. The table below describes the supported types and how read and write commands are handled for each reservation type.

Table 169: PERSISTENT RESERVE OUT, Type Code

Code	Name	Description
0h	Reserved	Reserved
1h	Write Exclusive	Reads Shared: Any initiator may execute commands that transfer from the media. Writes Exclusive: Only the initiator with the reservation may execute commands that transfer data to the media; Reservation Conflict status will be returned to other initiators.
2h	Reserved	Reserved
3h	Exclusive Access	Reads Exclusive: Only the initiator with the reservation may execute commands that transfer data from the media; Reservation Conflict status will be returned to other initiators. Writes Exclusive: Only the initiator with the reservation may execute commands that transfer data to the media; Reservation Conflict status will be returned to other initiators.

4h	Reserved	Reserved
5h	Write Exclusive Registrants Only	Reads Shard: Any initiator may execute commands that transfer from media. Writes Exclusive: Only registered initiators may execute commands that transfer data to the media; Reservation Conflict status will be returned to other initiators.
6h	Exclusive Access Registrants Only	Reads Exclusive: Only registered initiators may execute commands that transfer data from the media; Reservation Conflict status will be returned to other initiators. Writes Exclusive: Only registered initiators may execute commands that transfer data to the media; Reservation Conflict status will be returned to other initiators.
7h-Fh	Reserved	Reserved

17.13.3 Parameter list

The **Parameter List** required to perform the PERSISTENT RESERVE OUT command is defined in the table below. All fields must be sent on all PERSISTENT RESERVE OUT commands, even if the field is not required for the specified service action.

Table 170: Parameter List

Byte	Bit							
	7	6	5	4	3	2	1	0
0-7	(MSB) Reservation Key							(LSB)
8-15	(MSB) Service Action Reservation Key							(LSB)
16-19	(MSB) Reserved = 0							(LSB)
20	Reserved = 0			SPEC_I_P	ALL_TG_P T	Reserved = 0		APTPL
21-23	Reserved = 0							

Reservation Key contains an 8-byte value provided by the initiator, and identifies the initiator that issued the PERSISTENT RESERVE OUT command. The Reservation Key must match the registered reservation key for the initiator for all service actions except REGISTER and REGISTER AND IGNORE EXISTING KEY.

Service Action Reservation Key contents vary based on the service action. For REGISTER and REGISTER AND IGNORE EXISTING KEY, the Service Action Reservation Key must contain the new reservation key to be registered. For PREEMPT and PREEMPT AND ABORT, the field contains the reservation key of the persistent reservation that is being preempted. This field is ignored for all other service actions.

If the Specify Initiator Ports (**SPEC_I_PT**) bit is set to zero, the device server shall apply the registration only to the I_T nexus that sent the PERSISTENT RESERVE OUT command. If the SPEC_I_PT bit is set to one for any service action except the REGISTER service action, then the command shall be terminated 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 SPEC_I_PT bit is set to one for the REGISTER service action, the additional parameter data (see table XXX) shall include a list of transport IDs and the device server shall also apply the registration to the I_T nexus for each initiator port specified by a Transport ID. If a registration fails for any initiator port (e.g., if the logical unit does not have enough resources available to hold the registration information), no registrations shall be made, and the command shall be terminated with CHECK CONDITION status.

For Transport IDs, please refer to Table 179 on page 306 of Spc4r27.

The All Target Ports (**ALL_TG_PT**) bit is valid only for the REGISTER service action and the REGISTER AND IGNORE EXISTING KEY service action, and shall be ignored for all other service actions. Support for the ALL_TG_PT bit is optional. If the device server receives a REGISTER service action or a REGISTER AND IGNORE EXISTING KEY service action with the ALL_TG_PT bit set to one, it shall create the specified registration on all target ports in the SCSI target device known to the device server (i.e., as if the same registration request had been received individually through each target port). If the device server receives a REGISTER service action or a REGISTER AND IGNORE EXISTING KEY service action with the ALL_TG_PT bit set to zero, it shall apply the registration only to the target port through which the PERSISTENT RESERVE OUT command was received.

APTPL (Activate Persist Through Power Loss) bit is valid only for REGISTER and REGISTER AND IGNORE EXISTING KEY, and is ignored for all other service actions. If the last valid APTPL bit value received is zero, power loss will cause all persistent reservations to be released, and all reservation keys to be removed. If the last valid APTPL bit value received is one, any persistent reservation and all reservation keys for all initiators will be retained across power cycles.

17.13.4 Summary

Table 171: PERSISTENT RESERVE OUT, Service Action, Parameters

Service Action	Parameters						Generation counter
	Scope Type	Rsv Key	SvcAct RsvKey	S-spec addr	Extent length	APTPL	
(0) Register	ignore	verify	save	ignore	ignore	apply	+ 1
(1) Reserve	apply	verify	ignore	ignore	ignore	ignore	---
(2) Release	apply	verify	ignore	ignore	ignore	ignore	---
(5) Preempt and Abort	apply	verify	save	ignore	ignore	ignore	+ 1

17.13.4.1 Scope, Type

The Scope and the Type are applied in the process for the Reserve, Release, and Preempted and Clear service action but they are ignored in the process for the Register service action because they are not used.

17.13.4.2 Reservation Key

The Reservation Key is verified in each service action process. If the Initiator that registered a key is different from the Initiator requesting PERSISTENT RESERVE OUT command, the drive returns a **Reservation Conflict** status.

17.13.4.3 Service Action Reservation Key

On Register service action, the drive saves the key specified in the Service Action Reservation Key field as a key of Initiator requesting PERSISTENT RESERVE OUT command.

On Preempt and Clear service action, the reservation that has a key specified in the Service Action Reservation Key field is preempted.

On other service actions, this field is ignored.

17.13.4.4 APTPL

The APTPL (Active Persist Through Power Loss) is valid only for the Register service action. The drive ignores the APTPL in other service actions.

The following table shows the relationship between the last valid APTPL value and information held by the drive.

Table 172: APTPL and information held by a drive

Information held by the drive	The last valid APTPL value	
	0	1
Registration	all keys are set to 0	retained
Persistent Reservation	all are removed	retained
Generation counter	set to 0	set to 0

17.13.4.5 Generation counter

The drive increments the Generation counter when Register service action or Preempt and Clear service action complete successfully.

17.14 PRE-FETCH (34)

Table 173: PRE-FETCH (34)

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Command Code = 34h							
1	Reserved = 0			Reserved = 0			Immed = 0	Obsolete
2-5	(MSB) Logical Block Address (LSB)							
6	Reserved = 0							
7-8	(MSB) Transfer Length (LSB)							
9	VU = 0	Reserved = 0				FLAG	LINK	

The PRE-FETCH command requests the drive to transfer data to the cache. This command is implemented as a no-op and returns good status on the SSD.

17.15 READ (6) - (08)

Table 174: READ (6) - (08)

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Command Code = 08h							
1	Reserved = 0			(MSB) LBA				
2-3	Logical Block Address (LSB)							
4	Transfer Length							
5	VU = 0		Reserved = 0				FLAG	LINK

The READ command requests the drive to transfer from the medium to the initiator the specified number of blocks (Transfer Length) starting at the specified Logical Block Address (LBA).

- **Logical block address** field specifies the logical unit at which the READ operation shall begin.
- **Transfer length** field specifies the number of blocks to be transferred. A value of zero implies 256 blocks are to be transferred.

17.16 READ (10) - (28)

Table 175: READ (10) - (28)

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Command Code = 28h							
1	RDPROTECT		DPO	FUA	Rsvd=0	FUA_NV	Obsolete	
2-5	(MSB) Logical Block Address							(LSB)
6	Reserved = 0							
7-8	(MSB) Transfer Length							(LSB)
9	VU = 0	Reserved = 0				FLAG	LINK	

The READ (10) command requests the drive to transfer data to the Initiator. The larger LBA and Transfer Length fields permit greater quantities of data to be requested per command than with the READ command and are required to access the full LBA range of the larger capacity drives.

- **FUA_NV** (Force Unit Access Non-Volatile Cache) may be set to 0 or 1, but is ignored since NV_SUP=0 in Inquiry Page 86h.
- **Transfer length** The number of contiguous blocks to be transferred. If the transfer length is zero, the seek occurs, but no data is transferred. This condition is not considered an error. If read ahead is enabled, a read ahead is started after the seek completes.
- **DPO** (Disable Page Out) bit is ignored.
- **FUA** (Force Unit Access) bit is ignored.
- **RDPROTECT** defines the manner in which protection information read from drive shall be checked during processing of the command. Protection information is stored on drive, and may be transmitted to the drive's internal data buffer and to the initiator with the user data. If the drive is not formatted with protection information, RDPROTECT must be set to 000b, else **Check Condition** status will be returned with sense key of Illegal Request and additional sense code of Invalid Field in CDB.
- RDPROTECT=000b

Protection information is not transmitted to the initiator and is not checked.

RDPROTECT=001b

- Protection information is transmitted to the initiator with the user data
- Logical Block Guard is checked
- Logical Block Application Tag is checked (applies to READ(32) command only)
- Logical Block Reference Tag is checked

RDPROTECT=010b

- Protection information is transmitted to the initiator with the user data
- Logical Block Guard is not checked
- Logical Block Application Tag is checked (applies to READ(32) command only)
- Logical Block Reference Tag is checked

RDPROTECT=011b

- Protection information is transmitted to the initiator with the user data
- Logical Block Guard is not checked
- Logical Block Application Tag is not checked
- Logical Block Reference Tag is not checked

RDPROTECT=100b

- Protection information is transmitted to the initiator with the user data
- Logical Block Guard is checked
- Logical Block Application Tag is not checked
- Logical Block Reference Tag is not checked

RDPROTECT=101b, 110b, 111b

These values are reserved. **Check Condition** status will be returned with sense key of Illegal Request and additional sense code of Invalid Field in CDB.

If a check of the protection information fails, **Check Condition** status will be returned with sense key of Aborted Command and additional sense code indicating which protection field check failed.

Refer to the ANSI T10 standards for additional details of protection information.

If the transfer length is zero, no data is transferred. The CDB is validated and protocol checked and, if no problems are found, **Good** status is returned immediately. This condition is not considered an error.

17.17 READ (12) - (A8)

Table 176: Read (12) - (A8)

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Command Code = A8h							
1	RDPROTECT		DPO	FUA	Rsvd = 0	FUA_NV	Rsvd = 0	
2 - 5	(MSB) Logical Block Address (LSB)							
6 - 9	(MSB) Transfer Length (LSB)							
10	Reserved = 0							
11	VU = 0		Reserved = 0			FLAG	LINK	

The READ(12) command causes the drive to transfer data to the initiator. See the READ(10) description for the definitions of the fields in this command.

17.18 READ (16) - (88)

Table 177: READ (16) - (88)

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Command Code = 88h							
1	RDPROTECT			DPO	FUA	Rsvd=0	FUA_N V	Rsvd=0
2 - 9	(MSB) Logical Block Address (LSB)							
10-13	(MSB) Transfer Length (LSB)							
14	Restricted For MMC-4	Reserved = 0		GROUP NUMBER				
15	VU = 0		Reserved = 0			FLAG	LINK	

The READ(16) command causes the drive to transfer data to the initiator. See the READ(10) description for the definitions of the fields in this command.

17.19 READ (32) - (7F/09)

Table 178: READ (32) - (7F/09)

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Command Code = 7Fh							
1	VU = 0		Reserved = 0				FLAG	LINK
2-5	Reserved = 0							
6	Reserved = 0			Group Number = 0				
7	Additional CDB Length = 18h							
8 - 9	Service Action = 0009h							
10	RDPROTECT			DPO	FUA	Rsvd=0	FUA_NV	Rsvd=0
11	Reserved = 0							
12 -19	(MSB) Logical Block Address (LSB)							
20 - 23	(MSB) Expected Initial Logical Block Reference Tag (LSB)							
24 - 25	(MSB) Logical Block Application Tag (LSB)							
26-27	(MSB) Logical Block Application Tag Mask (LSB)							
28 - 31	(MSB) Transfer Length (LSB)							

The READ command requests that the drive transfer data from drive to the initiator. Each logical block transferred includes user data and may include protection information, based on the RDPROTECT field and the drive format.

If the drive is formatted with type 2 protection (PROT_EN=1 and P_TYPE=001b in the READ CAPACITY (16) parameter data), then this command will be processed normally. Any other protection types will result in Check Condition status to be returned with sense key of Illegal Request and additional sense code of Invalid Command Operation Code

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.

If the ATO bit is set to one in Mode Page 0Ah, 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 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 the protection information.

If the ATO bit is set to zero, the Logical Block Application Tag Mask field and the Expected Logical Block Application Tag field are ignored.

17.20 READ BUFFER (3C)

Table 179: READ BUFFER (3C)

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Command Code = 3Ch							
1	Reserved = 0			Mode				
2	Buffer ID = 0							
3-5	(MSB) Buffer Offset							(LSB)
6-8	(MSB) Allocation Length							(LSB)
9	VU = 0		Reserved = 0			FLAG	LINK	

The READ BUFFER command is used in conjunction with the WRITE BUFFER command as a diagnostic function for testing the memory of the drive and the SCSI bus integrity. This command does not alter the medium.

The function of this command and the meaning of fields within the command descriptor block depend on the contents of the mode field.

MODE	Description
00000	Read Combined Header and Data
00010	Read Data
00011	Descriptor
01010	Read Data from Echo Buffer
01011	Echo Buffer Descriptor
11010	Enable Expander Communications Protocol and Echo Buffer
All others	Not supported

17.20.1 Combined Header And Data (Mode 00000b)

In this mode a 4-byte header followed by data bytes is returned to the Initiator during the DATA IN phase. The Buffer ID and the buffer offset field are reserved.

The drive terminates the DATA IN phase when allocation length bytes of header plus data have been transferred or when the header and all available data have been transferred to the Initiator, whichever is less.

The 4-byte READ BUFFER header (see figure below) is followed by data bytes from the data buffer of the drive.

Table 180: Read Buffer Header

Byte	Bit						
	7	6	5	4	3	2	1
0	RSVD = 0						
1-3	(MSB) Buffer Capacity (LSB)						

The buffer capacity specifies the total number of data bytes that are available in the data buffer of the drive. This number is not reduced to reflect the allocation length nor is it reduced to reflect the actual number of bytes written using the WRITE BUFFER command.

Following the READ BUFFER header the drive will transfer data from its data buffer.

17.20.2 Read Data (Mode 00010b)

In this mode, the DATA IN phase contains buffer data.

- **Buffer ID** field must be set to zero, indicating the data transfer buffer. If another value is specified, the command is terminated with **Check Condition** status. The drive shall set sense key to *Illegal Request* and additional sense code to *Illegal Field in CDB*.
- **Buffer Offset** specifies the offset of the memory space specified by the Buffer ID. The Initiator should conform to the offset boundary requirements returned in the READ BUFFER descriptor. If the value exceeds the buffer specified, the command is terminated with **Check Condition** status. The drive shall set sense key to *Illegal Request* and additional sense code to *Illegal Field in CDB*.
- **Allocation Length** The drive terminates the DATA IN phase when allocation length bytes of data have been transferred or when the header and all available data have been transferred to the Initiator, whichever is less.

17.20.3 Descriptor (Mode 00011b)

In this mode, a maximum of four bytes of READ BUFFER descriptor information are returned. The drive returns the descriptor information for the buffer specified by the Buffer ID.

- **Buffer ID** field should normally be set to zero, indicating the drive data transfer buffer. If any other value is specified, the drive returns all zeros in the READ BUFFER descriptor.
- **Buffer Offset** field is reserved.
- **Allocation Length** should be set to four or greater. The drive transfers the allocation length or four bytes of READ BUFFER descriptor, whichever is less. The allocation length of zero indicates no data is transferred. The allocation length of greater than zero and less than four (size of the Descriptor) is an invalid request and will cause the command to be terminated with **Check Condition** status. The drive shall set sense key to *Illegal Request* and additional sense code to *Illegal Field in CDB*.

The READ BUFFER descriptor is defined in the figure below.

Table 181: Read Buffer Description

Byte	Bit						
	7	6	5	4	3	2	1
0	Offset Boundary = 0x09						
1-3	(MSB) Buffer Capacity (LSB)						

The value contained in the Buffer Offset field of subsequent WRITE BUFFER and READ BUFFER commands should be a multiple of two to the power of the offset boundary. The offset boundary is always set to nine, which indicates Sector Boundaries.

17.20.4 Read Data from Echo Buffer (Mode 01010b)

In this mode the drive transfers data from the echo buffer. The echo buffer will transfer the same data as when the WRITE BUFFER command was issued with the mode field set to echo buffer.

WRITE BUFFER command with the mode field set to echo buffer should be sent prior to the READ BUFFER command; otherwise the READ BUFFER command will be terminated with **Check Condition** status and *Illegal Request*.

In this mode Read Buffer transfers the specified amount of data or the amount previously written with a Write Buffer using mode 1010b from the echo buffer, whichever is less.

Issuing a Read Buffer mode 1010b before a Write Buffer mode 1010b will cause indeterminate data to be returned.

The most significant two bytes of the Allocation Length are ignored. The specified amount of data transferred should not be larger than the echo buffer capacity. The echo buffer capacity may be determined by using Read Buffer mode 1011b. Any additional data transferred over and above the echo buffer capacity is regarded as indeterminate.

The Buffer ID and Buffer Offset fields are ignored in this mode.

Note: The echo buffer is a separate buffer from the data buffer used with other read buffer modes. It is intended to be used for domain validation purposes.

17.20.5 Echo Buffer Descriptor (Mode 01011b)

In this mode, a maximum of four bytes of Read Buffer Descriptor information is returned. The drive returns the descriptor information for the echo buffer. The Buffer Offset field is reserved in this mode and must be zero. The drive transfers the lesser of the allocation length or four bytes of following Echo Buffer Descriptor.

Table 182: Echo Buffer Descriptor

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Reserved = 0							EBOS=0
1	Reserved = 0							
2	Reserved = 0				(MSB) Buffer Capacity			
3	Buffer Capacity (LSB)							

- **EBOS** (Echo Buffer Overwritten Supported) bit of zero indicates that the echo buffer is shared by all Initiators.
- **Buffer Capacity** field returns the size of the echo buffer in byte aligned to a 4-byte boundary.

17.20.6 Expander Communications and Echo Buffer (Mode 11010b)

Receipt of a READ BUFFER command with this mode (11010b) causes a communicative expander to enter the expanded communication protocol mode. SCSI target devices that receive a READ BUFFER command with this mode shall process it as if it were a READ BUFFER command with mode 01010b (see 17.17.4 Read Data from Echo Buffer).

17.21 READ CAPACITY (10) - (25)

Table 183: READ CAPACITY (10) - (25)

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Command Code = 25h							
1	Reserved = 0			Reserved = 0				Obso- lete
2-5	(MSB) Logical Block Address							(LSB)
6-7	Reserved = 0							
8	Reserved = 0						PMI	
9	VU = 0		Reserved = 0			FLAG	LINK	

The READ CAPACITY command returns information regarding the capacity of the drive.

- **Logical Block Address** is used in conjunction with the PMI bit.
- **PMI** (Partial Medium Indicator) indicates:

PMI Description

- | | |
|----------|--|
| 0 | The drive returns the last LBA of the drive. If the LBA field is not 0, the command to be terminated with Check Condition status. The drive shall set sense key to Illegal Request and additional sense code to Illegal Field in CDB. |
| 1 | The drive returns the last LBA and block length in bytes are that of the LBA after which a substantial delay in data transfer will be encountered. On SSDs, this returned LBA is the last LBA of the drive. If the LBA field is greater than the last LBA of the drive, the command to be terminated with Check Condition status. The drive shall set sense key to Illegal Request and additional sense code to Illegal Field in CDB |

17.21.0.1 Returned Data Format

The data returned to the Initiator in response to the READ CAPACITY command is described here. The data is returned in the DATA IN phase.

Table 184: Format of READ CAPACITY command reply

Byte	Bit							
	6	7	5	4	3	2	1	0
0-3	(MSB) Maximum Logical Block Address (LSB)							
4-7	(MSB) Block Length (LSB)							

- **Block Length** specifies the length in bytes of each block of user data (not including protection information).

17.22 READ CAPACITY (16) (9E/10)

Table 185: Read Capacity (16) (9E/10)

Byte	Bit							
	6	7	5	4	3	2	1	0
0	Command Code = 9Eh							
1	Reserved = 0				Service Action = 10h			
2-9	(MSB) Logical Block Address							(LSB)
10-13	(MSB) Allocation Length							(LSB)
14	Reserved = 0							PMI
15	CONTROL							

The READ CAPACITY (16) (9E/10) command returns information regarding the capacity of the drive. This command is processed like the standard READ CAPACITY (25) command. The contents of the CONTROL byte are defined in SAM-4.

17.22.1 Returned Data Format.

Table 186: Returned Data Format

Byte	Bit							
	6	7	5	4	3	2	1	0
0 - 7	(MSB) Returned Logical Block Address							(LSB)
8 - 11	(MSB) Logical Block Length in Bytes							(LSB)
12	Reserved = 0				P_Type		Prot_En	
13	P_I_Exponent				Logical Block Per Physical Block Exponent			
14	TPE	TPRZ	(MSB) Lowest Aligned Logical Block Address					
15	Lowest Aligned Logical Block Address							(LSB)
16 - 31	Reserved							

The RETURNED LOGICAL BLOCK ADDRESS field and LOGICAL BLOCK LENGTH IN BYTES field of the READ CAPACITY (16) parameter data are the same as the in the READ CAPACITY (10) parameter data.

The maximum value that shall be returned in the RETURNED LOGICAL BLOCK ADDRESS field is FFFF_FFFF_FFFF_FFEh.

The protection type (P_TYPE) field and the protection enable (PROT_EN) bit indicate the logical unit's current type of protection. See table below:

Table 187: P_TYPE field and PROT_EN bit

PROT_EN	P_TYPE	Description
0	xxx b	The logical unit is formatted to type 0 protection
1	000 b	The logical unit is formatted to type 1 protection
1	001 b	The logical unit is formatted to type 2 protection
1	010 b	The logical unit is formatted to type 3 protection
1	011 b to 111 b	Reserved

The **P_I_EXPONENT** field may be used to determine the number of protection information intervals placed within each logical block.

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 Block Per Physical Block Exponent** field is defined below:

Table 188: 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 Equation
1	The logical unit is formatted to type 2 protection
(a) The number of physical blocks per logical block is not reported.	

If the thin provisioning enabled (**TPE**) bit is set to one, then the logical unit implements thin provisioning
 If the TPE bit is set to zero, then the logical unit implements full provisioning

If the thin provisioning read zeros (**TPRZ**) bit is set to one, then, for an unmapped LBA specified by a read operation, the device server shall send user data with all bits set to zero to the data-in buffer. If the TPRZ bit is set to zero, then, for an unmapped LBA specified by a read operation, the device server shall send user data with all bits set to any value to the data-in buffer.

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.

NOTE: The highest LBA that the lowest aligned logical block address field supports is 3FFFh (i.e.,16,383).

17.23 READ DEFECT DATA (37)

Table 189: READ DEFECT DATA (37)

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Command Code = 37h							
1	Reserved = 0			Reserved = 0				0
2	Reserved = 0			Plist	Glist	Defect List Format		
3-6	Reserved = 0							
7-8	(MSB) Allocation Length							(LSB)
9	VU = 0		Reserved = 0			FLAG	LINK	

The READ DEFECT DATA command requests that the Target transfer the medium defect data to the Initiator.

If the Target is unable to access any medium defect data it will return a **Check Condition** status with the appropriate sense key. The sense key will be set to either *Medium Error* (03h) if a medium error occurred or *No Sense* (00h) if the list does not exist and the additional sense code will be set to *Defect List Error* (19h).

- **Plist** bit set to one indicates that the Target returns the Plist. A Plist bit of zero indicates that the Target shall not return the Plist of defects.
- **Glist** bit set to one indicates that the Target returns the Glist. A Glist bit of zero indicates that the Target shall not return the Glist.

Note: With both bits set to one Plist and Glist the Target will return both the primary and grown defect lists. With both bits set to zero, the Target will return only a 4-byte Defect List Header.

- **Defect List format** field is used by the Initiator to indicate the preferred format for the defect list.

The Defect List Format of '101 (Physical Sector Format)' is supported. If the requested format is not supported by the drive, it will return the defect list in its default format '101' and then terminate the command with **Check Condition** status. The sense key will be set to *Recovered Error* (01h) and the additional sense code will be set to *Defect List Not Found* (1Ch).

The drive sends defect list (Defect Descriptors) in a 8-byte Absolute Block Address (ABA) format that follows a four byte Defect List Header.

The Target will transfer all of the Read Defect Data up to the number of bytes allocated by the Initiator.

Table 190: Defect List Format

Preferred Defect List Format	Returned Defect List Format
Block (000b)	Physical Sector
Bytes from Index (100b)	Physical Sector
Physical Sector (101b)	Physical Sector
Vendor Unique (110b)	Physical Sector
Reserved (001b)	
Reserved (010b)	
Reserved (011b)	
Reserved (111b)	

Note: The drive will terminate the Data In phase when the Allocation Length has been transferred or when all available Defect Data has been transferred to the Initiator, whichever is less.

The Read Defect Data contains a 4-byte header followed by zero or more defect descriptors.

17.23.1 Defect List Header

Table 191: Defect List Header

Byte	Bit							
	7	6	5	4	3	2	1	0
	Defect List Header							
0	Reserved = 0							
1	Reserved = 0			Plist	Glist	Defect List Format		
2-3	(MSB)			Defect List length				(LSB)

17.23.2 Defect List Descriptor

Table 192: Defect List Descriptor

Byte	Bit							
	7	6	5	4	3	2	1	0
	Defect List Descriptor							
0-7	Defect Descriptor 0							
.								
8n - (8n+7)	Defect Descriptor n							

17.23.3 Physical Sector Format (101b)

Table 193: Defect Descriptors of Physical Sector Format

Byte	Defect Descriptors
3-0	(MSB) Die of Defect (LSB)
7-4	(MSB) Erase Block of Defect (LSB)

The Defect List Format field specifies the format of the defect list data returned by the Target.

The Defect List Length field specifies the length in bytes of the defect descriptors that follow. The Defect List Length is equal

to eight times the number of defect descriptors.

Normally the Target will set the Defect List Length field to the amount of space needed to contain the entire defect list. However, the Target is capable of building a defect list with a length such that the entire list cannot be transferred using the maximum allocation length. If the defect list grows beyond 8191 entries, the defect data cannot be transferred with an allocation length of 0FFFFh. The Target will transfer a partial defect list and return **Check Condition** status with the sense key set to *Recovered Error* and the additional sense code set to *Partial Defect List Transferred*. The defect list length will be set to 0FFF8h, indicating the maximum number of defect descriptors that can be transferred. Defects beyond this number cannot be read by the Initiator.

17.24 READ DEFECT DATA (B7)

Table 194: READ DEFECT DATA (B7)

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Command Code = B7h							
1	Reserved = 0			Plist	Glist	Defect List Format		
2-5	Reserved = 0							
6-9	(MSB) Allocation Length							(LSB)
10	Reserved = 0							
11	VU = 0		Reserved = 0			FLAG	LINK	

(See Section 17.23 READ DEFECT DATA (37)” on page 203.)

17.24.1 Defect List Header

Table 195: Defect List Header

Byte	Bit							
	7	6	5	4	3	2	1	0
	Defect List Header							
0	Reserved = 0							
1	Reserved = 0			Plist	Glist	Defect List Format		
2-3	Reserved = 0							
4-7	(MSB) Defect List length							(LSB)

(See Defect List Header for Read Defect Data (37) in Section Table 17.23.1, “Defect List Header,” on page 205.)

17.24.2 Defect List Descriptor

Table 196: Defect List Descriptor

Byte	Bit							
	7	6	5	4	3	2	1	0
	Defect List Descriptor							
0-7	Defect Descriptor 0							
.								
8n - (8n+7)	Defect Descriptor n							

(See Defect List Descriptor for Read Defect Data (37) in Section 17.23.2 Defect List Descriptor” on page 205.)

17.24.3 Physical Sector Format (101b)

Table 197: Defect Descriptors of Physical Sector Format

Byte	Defect Descriptors
0-3	(MSB) Die of Defect (LSB)
4-7	(MSB) Erase Block of Defect (LSB)

17.25 READ LONG (3E)

Table 198: READ LONG (3E)

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Command Code = 3Eh							
1	Reserved = 0			Reserved = 0			Correct = 0	Obsolete
2-5	(MSB) Logical Block Address (LSB)							
6	Reserved = 0							
7-8	(MSB) Byte Transfer Length (LSB)							
9	VU = 0	Reserved = 0				FLAG	LINK	

The READ LONG command requests the drive to transfer one block of data to the Initiator. The transfer data includes data and ECC field data.

- **Correct** bit is ignored. ECC correction is always performed. If ECC correction fails, the Target terminates the command with Check Condition status, the sense key is set to Medium Error, and an additional sense code set to Unrecovered Read Error.
- **Logical Block Address** field specifies the logical block at which the read operation shall occur.
- **Byte Transfer Length** field must specify exactly the number of bytes of data that are available for transfer. If a non-zero byte transfer length does not match the available data length, the Target terminates the command with **Check Condition** status, the sense key is set to *Illegal Request*, and an additional sense code set to *Invalid Field in CDB*. The valid and ILI bits are set to one and the information field is set to the difference of the requested length minus the actual length in bytes. Negative values are indicated by two's complement notation.

The transfer length is calculated as follows:

$$transfer\ length = logical\ block\ size + 4$$

The data read by this command is neither read from nor retained in the cache.

17.26 REASSIGN BLOCKS (07)

Table 199: REASSIGN BLOCKS (07)

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Command Code = 07h							
1	Reserved = 0				Reserved = 0			
2	Reserved = 0							
3								
4								
5	VU = 0	Reserved = 0				FLAG	LINK	

The REASSIGN BLOCKS command is implemented as a no-op on the SSD. To maintain compatibility, the SSD performs protocol checking on the CDB and a range check of the LBA(s) transferred to the drive during the DATA OUT phase.

The REASSIGN BLOCKS command will not modify the specified LBAs or attempt to recover or reallocate them. An unreadable LBA will remain unreadable after execution of a REASSIGN BLOCKS command.

Following is the format of the data sent by the Initiator during the DATA OUT phase.

Table 200: Format of Reassign Blocks data

Byte	Bit						
	7	6	5	4	3	2	1
0	Reserved = 0						
1	Reserved = 0						
2-3	(MSB) Defect List Length = 4/8/12/16 (LSB)						
4-7	(MSB) Defect Logical Block Address 1 (LSB)						
8-11	(MSB) Defect Logical Block Address 2 (LSB)						
12-15	(MSB) Defect Logical Block Address 3 (LSB)						
16-19	(MSB) Defect Logical Block Address 4 (LSB)						

- **Defect List Length** must be 4, 8, 12, or 16. Otherwise, the drive returns *Check Condition* with a sense key of *Illegal Request*.
- **Defective Logical Block Address** is four bytes in length. The Initiator can specify from 1 to 4 Defective Logical Block Addresses according to the Defect List Length from 4 to 16, respectively. LBAs are not required to be in ascending order. If the Defective Logical Block Address is greater than the maximum LBA of the drive, the command will be terminated with *Check Condition* with a sense key of *Illegal Request*,

17.27 RECEIVE DIAGNOSTICS RESULTS (1C)

Table 201: RECEIVE DIAGNOSTIC RESULTS (1C)

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Command Code = 1Ch							
1	Reserved = 0			Reserved = 0			PCV	
2	Page Code							
3	(MSB) Allocation Length							
4					(LSB)			
5	VU = 0		Reserved = 0			FLAG		LINK

The RECEIVE DIAGNOSTIC RESULTS command requests that analysis data requested by a SEND DIAGNOSTIC command be sent to the Initiator.

- **PCV** (Page Code Valid) bit of zero indicates that the most recent SEND DIAGNOSTIC command shall define the data returned by this command. PCV bit of one indicates that the contents of the Page Code field shall define the data returned by this command.
- **Allocation Length** specifies the amount of data to be returned to the Initiator. This value may be zero and this is not considered an error. The Target terminates the Data In phase when all available data has been transferred or when the number of bytes transferred equals the Parameter List Length.

17.27.1 Receive Diagnostic Results Page 0

This page contains a list of supported pages.

Table 202: Receive Diagnostic Results page 0

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Page Code = 0							
1	Reserved = 0							
2-3	Page Length = 12h							
4	(Supported Pages) Page = 0h							
5-19	ESI Pages = 01h - 0Fh							
20	Translate address page = 40h							
21	Device LED Control Page = A0h							

The supported diagnostic page returns a list of supported pages in ascending order.

17.27.2 Enclosure Service Information (ESI) Page Format

The drive supports the following enclosure pages as specified by the “SCSI-3 Enclosure Service (SES) Rev 8a” standard. Please refer to that standard for more definition on these pages. Note that the drive does not attempt to process the information in these pages, but only acts as a pass through node, to allow the initiator to communicate with the enclosure.

Table 203: Enclosure Page Support for Send and Receive Diagnostic Commands

Page Code	Send Diagnostic Command	Receive Diagnostic Command
01h	Reserved	Configuration
02h	Enclosure Control	Enclosure Status
03h	Reserved	Help Text
04h	String Out	String In
05h	Threshold Out	Threshold In
06h	Array Control	Array Status
07h	Reserved	Element Descriptor
08h	Reserved	Short Enclosure Status
09-0Fh	Reserved for SES	Reserved for SES

These pages are formatted as in the following table.

Table 204: ESI Page Format

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Page Code = 1h-0Fh							
1	Page Specific							
2	(MSB) Page Length (LSB)							
3								
n	Page Specific							

If the enclosure supports 'short' mode, only page 08h is supported. If the enclosure supports 'long' mode, pages 01h-0Fh are supported.

17.27.3 Receive Diagnostic Results Page 40

Using the SEND DIAGNOSTIC command, an address in either physical or logical format is supplied to the drive. This page is then used to retrieve the address translated into the other format. Note that this address translation functionality is only supported by the SSD for HDD compatibility. The physical address does not represent a true physical address on the SSD's media.

Table 205: Receive Diagnostic Results Page 40

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Page Code = 40h							
1	Reserved = 0							
2-3	Page Length							
4	Reserved = 0					Supplied format		
5	RA	Reserved=0				Translate format		
6-n	Translated Address							

- **Page Length** is set to 02h if the address is in a Reserved Area (RA =1). Otherwise, Page Length is set to 06h if the Translate Format is Block format, or 0Ah if the Translate Format is Bytes From Index format or Physical Sector format.
- **Supplied Format** is the value supplied by the SEND DIAGNOSTIC command; it may be one of the three following values:
 - **000b** Block format
 - **100b** Bytes From Index format
 - **101b** Physical Sector format
- **Translate Format** is the value supplied by the SEND DIAGNOSTIC command and specifies the format in which the address has been translated into List. If the supplied format is the Block format, the Translate format must be either Bytes from Index or Physical Sector format. If the supplied format is the Bytes from Index or Physical Sector format, the Translate format must be Block format. Otherwise the Target will terminate the command with **Check Condition** status.
- **RA (Reserved Area)** is set to on if the translated block is an inaccessible sector, which reflects a defect, an unused sector on a spare cylinder, or a sector beyond the Maximum Customer LBA.
- **Translated Address** contains the address in the translate format. If it is an LBA, it is contained within the first four bytes of the field (bytes 6 to 9) of the page data. For a physical format it is as follows:

Table 206: Translated address

Byte	Bit							
	7	6	5	4	3	2	1	0
6-13	Physical Address							

17.27.4 Receive Diagnostic Page A0h

The Receive Diagnostic Device LED Control Page A0h is returned as a result of the Send Diagnostic Device LED Control Page A0h. For a description of the parameters in this page, see 19.34.3, “Send Diagnostic Page A0h” on page 231.

Table 207: Device LED Control Page - Receive Diagnostic

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Page Code = A0h							
1	Reserved = 0							
2	Page Length = 0002h							
3	(MSB)							(LSB)
4	Reserved = 0							
5	Reserved = 0							SBDL

17.28 RELEASE (17)

Table 208: RELEASE (17)

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Command Code = 17h							
1	Reserved = 0		3rdPty=0	3rd Party ID			Ext=0	
2	Reservation Identification							
3-4	Reserved = 0							
5	VU = 0		Reserved = 0			FLAG	LINK	

The RELEASE command is used to release a LUN previously reserved. It is not an error for an Initiator to release a LUN that is not currently active. The drive returns **Good** status without altering the reservation.

- **3rdPty** must be 0. Third Party reservations are not supported. If the 3rdPty bit is not zero, Check Condition status is returned with a sense key of Illegal Request and additional sense code of Invalid Field in CDB.
- **3rd Party ID** is ignored.
- **Extents** must be 0. Extension is not supported by the drive.
- **Reservation Identification** field is ignored.

17.29 RELEASE (57)

Table 209: RELEASE (57)

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Command Code = 57h							
1	Reserved = 0		3rdPty=0	Reserved = 0			Ext = 0	
2	Reservation Identification							
3	3rd Party Device ID							
4-8	Reserved = 0							
9	VU = 0		Reserved = 0			FLAG	LINK	

The RELEASE command is used to release a LUN previously reserved. It is not an error for an Initiator to release a LUN that is not currently active. The drive returns **Good** status without altering the reservation.

- **3rdPty** must be 0. Third Party reservations are not supported. If the 3rdPty bit is not zero, Check Condition status is returned with a sense key of Illegal Request and additional sense code of Invalid Field in CDB.
- **3rd Party ID** is ignored.
- **Extent** must be 0. Extension is not supported by the drive.
- **Reservation Identification** field is ignored.

17.30 REPORT DEVICE IDENTIFIER (A3/05)

Table 210: REPORT DEVICE IDENTIFIER (A3/05)

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Command Code = A3h							
1	Reserved = 0				Service Action = 05h			
2	Reserved = 0							
3	Reserved = 0							
4-5	(MSB) LUN=0 (LSB)							
6-9	(MSB) Allocation Length (LSB)							
10	Reserved = 0							
11	VU = 0		Reserved = 0				FLAG	LINK

The **REPORT DEVICE IDENTIFIER** command requests that the device server send device identification information to the application client.

The **LUN** contains the logical unit number parameter. This parameter is expected to be zero. Other value for this parameter will cause the command to terminate with a CHECK CONDITION status. The sense key is set to ILLEGAL REQUEST, and the additional sense code is set to INVALID FIELD IN CDB.

The **ALLOCATION LENGTH** field indicates how much space has been reserved for the returned parameter data. If the length is not sufficient to contain all the parameter data, the first portion of the data is returned. This is not considered an error. The actual length of the parameter data is available in the IDENTIFIER LENGTH field in the parameter data. If the remainder of the parameter data is required, the application client should send a new REPORT DEVICE IDENTIFIER command with an ALLOCATION LENGTH field large enough to contain all the data.

The REPORT DEVICE IDENTIFIER parameter list contains a 4-byte field that contains the length in bytes of the parameter list and the logical unit's identifier.

Table 211: Report Device Identifier parameter list

Byte	Bit							
	7	6	5	4	3	2	1	0
0-3	(MSB) Identifier Length = n - 3							(LSB)
4-n	Identifier							

The IDENTIFIER LENGTH field specifies the length in bytes of the IDENTIFIER field. If the ALLOCATION LENGTH field in the CDB is too small to transfer all of the identifier, the length is not adjusted to reflect the truncation. The identifier length initially equals zero and is changed only by a successful SET DEVICE IDENTIFIER command.

The IDENTIFIER field contains a vendor specific value. The value reported is the last value written by a successful SET DEVICE IDENTIFIER command. The value of the identifier is changed only by a SET DEVICE IDENTIFIER command. The identifier value persist through resets, power cycles, media format operations.

The Target return the same Identifier to all Initiators on all ports.

The execution of a REPORT DEVICE IDENTIFIER requires the enabling of a nonvolatile memory within the logical unit. If the nonvolatile memory is not ready, the device server returns **Check Condition** status rather than wait for the device to become ready. The sense key is set to *Not Ready* and the additional sense data is set as described in the TEST UNIT READY command. This information should allow the application client to determine the action required to cause the device server to become ready.

17.31 REPORT LUNS (A0)

Table 212: REPORT LUNS (A0)

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Command Code = A0h							
1-5	Reserved							
6-9	(MSB) Allocation Length (LSB)							
10	Reserved							
11	VU = 0		Reserved = 0				FLAG	LINK

The REPORT LUNS command requests that the Target return the known LUN to the Initiator. The REPORT LUNS command should always be available and is unaffected by any reservations.

The Allocation Length must be at least 16 bytes. If the Allocation Length is less than 16 bytes, the Target will return a **Check Condition** status with sense key of *Illegal Request* and additional sense code of *Invalid Field in CDB*. If the Allocation Length is not sufficient to contain the LUN values for all configured logical units, the Target shall report as many LUN values as will fit in the specified Allocation Length. This is not considered an error.

The REPORT LUNS command will send the LUN list in the subsequent Data Out Phase. The format of the LUN list is shown in the following table.

Table 213: LUN Reporting parameter list format

Byte	Bit							
	7	6	5	4	3	2	1	0
0-3	(MSB) LUN List Length = 8 (LSB)							
4-7	Reserved							
8-15	(MSB) LUN = 0 (LSB)							

The LUN list length shall contain the length in bytes of the LUN list that is available to be transferred. This product only supports one LUN. Therefore, the LUN list length must be set to 8. The only supported LUN is zero.

17.32 REPORT SUPPORTED OPERATION CODES (A3/0C)

Table 214: REPORT SUPPORTED OPERATION CODES (A3/0C)

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Command Code = A3h							
1	Reserved = 0				Service Action = 0Ch			
2	RCTD	Reserved = 0				Reporting Options		
3	Requested Operation Code							
4-5	Requested Service Action							
6-9	Allocation Length							
10	Reserved = 0							
11	VU = 0	Reserved				FLAG	LINK	

The REPORT SUPPORTED OPERATION CODES command requests information on commands that the drive supports. The initiator may request a list of all operation codes and service actions supported, or the command support data for a specific command.

RCTD: A return command timeouts descriptor (RCTD) bit set to one specifies that the command timeouts descriptor shall be included in each command descriptor (see section 17.32.1) that is returned or in the one_command parameter data (see section 17.32.2) that is returned. A RCTD bit set to zero specifies that the command timeouts descriptor shall not be included in any parameter data returned.

Reporting Options specifies the information to be returned in the parameter data.

Table 215: Reporting Options

Reporting Options	Description
000b	A list of all operation codes and service actions supported by the drive will be returned in the all_commands parameter data format. The Requested Operation Code field and Requested Service Action field will be ignored.
001b	The command support data for the operation code specified in the Requested Operation Code field will be returned in the one_command parameter data format. The Requested Service Action field will be ignored. If the Requested Operation Code field specifies an operation code that has service actions, Check Condition status will be reported with a sense key of Illegal Request and additional sense code of Invalid Field in CDB.
010b	The command support data for the operation code and service action specified in the Requested Operation Code field and Requested Service Action field will be returned in the one_command parameter data format. If the Requested Operation Code field specifies an operation code that does not have service actions, Check Condition status will be reported with a sense key of Illegal Request and additional sense code of Invalid Field in CDB.
011b-111b	Reserved

Requested Operation Code specifies the operation code of the command to be returned in the one_command parameter data format.

Requested Service Action specifies the service action of the command to be returned in the one_command parameter data format.

Allocation Length specifies the number of bytes that have been allocated for the returned parameter data. If the length is not sufficient to contain all the parameter data, the first portion of the data shall be returned. The actual length of the parameter data may be determined from the Additional Length field in the parameter data.

17.32.1 All_commands parameter data format

The Report Supported Operation Codes all_command parameter data format begins with a four-byte header that contains the length in bytes of the parameter data, followed by a list of supported commands. Each command descriptor contains information about a single supported command CDB (i.e. one operation code and service action combination, or one non-service action operation code).

Table 216: All_command parameter data format

Byte	Bit							
	7	6	5	4	3	2	1	0
0-3	Command Data Length (n-3)							
4	Command Descriptor 0							
N	Command Descriptor X							

Each **Command Descriptor** contains information about a single supported command CDB.

Table 217: Command Descriptor format

Byte	Bit							
	7	6	5	4	3	2	1	0
0-3	Operation Code							
1	Reserved = 0							
2-3	Service Action							
4	Reserved = 0							
5	Reserved = 0						CTDP	Serva- ctv
6-7	CDB Length							
8-19	Command Timeouts Descriptor, if any (see 17.32.3)							

Operation Code contains the operation code of a supported command.

Service Action contains a supported service action of the supported operation. If the operation code does not have a service

action, the Service Action field will be set to zero.

CTDP: A command timeouts descriptor present bit set to one indicates that the command timeouts descriptor (see 18.32.3) is included in this command descriptor. A CTDP bit set to zero indicates that the command timeouts descriptor is not included in this command descriptor.

Servactv set to zero indicates the operation code does not have service actions and the Service Action field should be ignored. **SERVACTV** set to one indicates the operation code field has service actions and the contents of the Service Action field are valid.

CDB Length contains the length of the command CDB in bytes.

17.32.2 One_command parameter data format

The Report Supported Operation Codes one_command parameter data format contains information about the CDB and a usage map for bits in the CDB for the command specified by the Reporting Options, Requested Operation Code, and Requested Service Action fields in the Reported Supported Operation Codes CDB.

Table 218: One_command parameter data format

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Reserved = 0							
1	CTDP	Reserved = 0				Support		
2-3	CDB Size (n-3)							
4-n	CDB Usage Data							
n+1 - n+12	Command Timeouts Descriptor, if any (see 17.32.3)							

CTDP: A command timeouts descriptor present bit set to one indicates that the command timeouts descriptor is included in this command descriptor. (see section 17.32.3) A CTDP bit set to zero indicates that the command timeouts descriptor is not included in this command descriptor.

The Support field is defined in the table below.

Table 219: One_command parameter support field

Recording Option	Description
000b	Data about the requested command is not currently available. All data after byte 1 is not valid. A subsequent request for command support data may be successful.
001b	The requested command is not supported. All data after byte 1 is not valid.
010b	Reserved.
011b	The requested command is supported in conformance with the standard.
100b	Reserved
101b	The requested command is supported in a vendor specific manner.
110b-111b	Reserved.

CDB Size contains the size of the CDB Usage Data field in the parameter data, and the number of bytes in the CDB for the command requested.

CDB Usage Data contains information about the CDB for the command requested. The first byte of the CDB Usage Data field contains the operation code for the command. If the command contains a service action, then that service action code is returned in the same location as the Service Action field of the command CDB. All other bytes of the CDB Usage Data field contain a usage map for bits in the CDB for the command requested.

The bits in the usage map have a one-for-one correspondence to the CDB for the command requested. If the drive evaluates a bit in the CDB, the usage map will contain a one in the corresponding bit position. The usage map will contain a zero in the corresponding bit position for any field treated as ignored or reserved.

17.32.3 Command timeouts descriptor format

17.32.3.1 Overview

The command timeouts descriptor (see Table 220:) returns time-out information for commands supported by the logical unit based on the time from the start of processing for the command to its reported completion. Values returned in the command timeouts descriptor do not include times that are outside the control of the device server (e.g., prior commands with the IMMED bit set to one in the CDB, concurrent commands from the same or different I_T nexuses, manual unloads, power-on self tests, prior aborted commands, commands that force cache synchronization, delays in the service delivery subsystem). For commands that cause a change in power condition (Idle/Standby Powersave Modes), values returned in the command timeouts descriptor do not include the power condition transition time (e.g., the time to spinup rotating media). Values returned in the command timeouts descriptor should not be used to compare products.

Table 220: Command timeouts descriptor format

Byte	7	6	5	4	3	2	1	0
0 - 1	Descriptor Length (0Ah)							
2	Reserved = 0							
3	Command Specific							
4- 7	Nominal Command Processing Time-out							
8 - 11	Recommended Command Time-out							

The **DESCRIPTOR LENGTH** field indicates the number of bytes that follow in the command timeouts descriptor.

The **COMMAND SPECIFIC** field contains time-out information (see Table 221:) that is specific to one or more commands. If no command specific time-out information is defined by this or the applicable command standard, the **COMMAND SPECIFIC** field is reserved.

Table 221: Command timeouts descriptor Command Specific Field usage

Command	Reference
WRITE BUFFER	See Section 17.32.3.2

17.32.3.2 WRITE BUFFER: command timeouts descriptor COMMAND SPECIFIC field usage

For the WRITE BUFFER command, the **COMMAND SPECIFIC** field usage is reserved for all modes except the following:

- Download microcode mode (04h);
- Download microcode and save mode (05h);
- Download microcode with offsets mode (06h);
- Download microcode with offsets and save mode (07h);
- Download microcode with offsets and defer activation mode (0Eh) only if the microcode is activated by an event other than an activate deferred microcode mode; and
- Activate deferred microcode mode (0Fh).

If the command timeouts descriptor describes one of the WRITE BUFFER modes listed in this subclause, then the **COMMAND SPECIFIC** field indicates the maximum time, in one second increments, that access to the SCSI device is limited or not possible through any SCSI ports associated with a logical unit that processes a WRITE BUFFER command that specifies one of the named modes. A value of zero in the **COMMAND SPECIFIC** field indicates that the no maximum time is indicated.

17.33 REPORT SUPPORTED TASK MANAGEMENT FUNCTIONS (A3/0D)

Table 222: Report Supported Task Management Functions (A3/0D)

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Command Code = A3h							
1	Reserved = 0			Service Action = 0Dh				
2-5	Reserved = 0							
6-9	Allocation Length							
10	Reserved = 0							
11	VU = 0		Reserved = 0				Flag	Link

The REPORT SUPPORTED TASK MANAGEMENT FUNCTIONS command requests information on task management functions supported by the drive.

Allocation Length specifies the number of bytes that have been allocated for the returned parameter data. The allocation length must be at least four. If the allocation length is less than four, Check Condition Status will be returned with sense key of Illegal Request and additional sense code of Invalid Field in CDB.

The format of the returned parameter data is shown below.

Byte	Bit							
	7	6	5	4	3	2	1	0
0	ATS	ATSS	CACAS	CTSS	LURS	QTS	TRS	WAKES
1-3	Reserved = 0							

ATS (Abort Task) bit set to one indicates that ABORT TASK is supported. An ATS bit of zero indicates that ABORT TASK is not supported.

ATSS (Abort Task Set) bit set to one indicates that ABORT TASK SET is supported. An ATSS bit of zero indicates that ABORT TASK SET is not supported.

CACAS (Clear ACA) bit set to one indicates that CLEAR ACA is supported. A CACAS bit of zero indicates that CLEAR ACA is not supported.

CTSS (Clear Task Set) bit set to one indicates that CLEAR TASK SET is supported. A CTSS bit of zero indicates that CLEAR TASK SET is not supported.

LURS (Logical Unit Reset) bit set to one indicates that LOGICAL UNIT RESET is supported. An LUR bit of zero indicates that LOGICAL UNIT RESET is not supported.

QTS (Query Task) bit set to one indicates that QUERY TASK is supported. A QTS bit of zero indicates that QUERY TASK is not supported.

TRS (Target Reset) bit set to one indicates that TARGET RESET is supported. A TRS bit of zero indicates that TARGET RESET is not supported.

WAKES (WakeUp) bit set to one indicates that WAKEUP is supported. A WAKES bit of zero indicates that WAKEUP is not supported.

17.34 REQUEST SENSE (03)

Table 223: REQUEST SENSE (03)

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Command Code = 03h							
1	Reserved = 0			Reserved = 0				
2-3	Reserved = 0							
4	Allocation Length							
5	VU = 0		Reserved = 0				FLAG	LINK

The REQUEST SENSE command requests the drive to transfer sense data.

If REQUEST SENSE command with an invalid LUN is received, the drive returns **Good** status and reports a sense key of *Illegal Request* and an additional sense code of *Logical Unit Not Supported*.

If the drive has no sense data available to return, it shall return a sense key of *No Sense* and an additional sense code of *No Additional Sense Information*.

Separate sense data is maintained by the device for each Initiator. Therefore, there is no requirement for an Initiator to expeditiously clear a *Check Condition* as this will not affect other initiators in a multi-Initiator system.

The drive will return the number of bytes in the allocation length or 32 bytes, whichever is less.

17.35 RESERVE (16)

Table 224: RESERVE (16)

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Command Code = 16h							
1	Reserved = 0		3rdPty=0	3rd Party ID			Ext=0	
2	Reservation Identification							
3-4	(MSB) Extent List Length = 0 (LSB)							
5	VU = 0		Reserved = 0			FLAG	LINK	

The RESERVE command is used to reserve a LUN for an Initiator. This reservation can be either for the Initiator sending the command or for a third party as specified by the Initiator.

Extents are not supported by the drive. The Ext bit must be zero. If Ext bit is set to one, **Check Condition** status is returned with a sense key of *Illegal Request* and additional sense code of *Invalid Field in CDB*. The Reservation Identification and Extent List Length fields are ignored.

The Reserve command requests that the entire LUN be reserved for the Initiator until

- the reservation is superseded by another valid Reserve command from the Initiator that made the reservation.
- the reservation is released by a RELEASE command from the same Initiator.
- a hard Reset condition occurs.
- a Target Reset message is received from any Initiator.
- a power off/on cycle occurs.

3rdPty must be 0. Third Party reservations are not supported. If the 3rdPty bit is not zero, Check Condition status is returned with a sense key of *Illegal Request* and additional sense code of *Invalid Field in CDB*.

3rd Party ID is ignored.

Only the Initiator that issued the Reserve command for a LUN may release the LUN, regardless of the 3rdPty option. This Initiator may also release the LUN by issuing another RESERVE command. This superseding RESERVE command releases the previous reservation when the new reservation is granted.

Reservation queuing is not supported by the drive. If a LUN is reserved and a RESERVE command is issued from a different Initiator, the Target responds with a RESERVATION CONFLICT.

17.36 RESERVE (56)

Table 225: RESERVE (56)

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Command Code = 56h							
1	Reserved = 0			3rdPty=0	Reserved			Ext=0
2	Reservation Identification							
3	Third Pay Device ID							
4-6	Reserved = 0							
7-8	(MSB) Extent List Length = 0							(LSB)
9	VU = 0		Reserved = 0			FLAG	LINK	

The RESERVE command is used to reserve a LUN for an Initiator. This reservation can be either for the Initiator sending the command or for a third party as specified by the Initiator.

Extents are not supported by the drive. The Ext bit must be zero. If Ext bit is set to one, **Check Condition** status is returned with a sense key of *Illegal Request* and additional sense code of *Invalid Field in CDB*. The Reservation Identification and Extent List Length fields are ignored.

The Reserve command requests that the entire LUN be reserved for the Initiator until

- the reservation is superseded by another valid Reserve command from the Initiator that made the reservation.
- the reservation is released by a RELEASE command from the same Initiator.
- a hard Reset condition occurs.
- a Target Reset message is received from any Initiator.
- a power off/on cycle occurs.

3rdPty must be 0. Third Party reservations are not supported. If the 3rdPty bit is not zero, Check Condition status is returned with a sense key of *Illegal Request* and additional sense code of *Invalid Field in CDB*.

3rd Party ID is ignored.

Only the Initiator that issued the Reserve command for a LUN may release the LUN, regardless of the 3rdPty option. This Initiator may also release the LUN by issuing another RESERVE command. This superseding RESERVE command releases the previous reservation when the new reservation is granted.

Reservation queuing is not supported by the drive. If a LUN is reserved and a RESERVE command is issued from a different Initiator, the Target responds with a RESERVATION CONFLICT.

17.37 REZERO UNIT (01)

Table 226: REZERO UNIT (01)

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Command Code = 01h							
1	Reserved = 0			Reserved = 0				
2-4	Reserved = 0							
5	VU = 0		Reserved = 0				FLAG	LINK

The REZERO UNIT command is implemented as a no-op on the SSD.

17.38 SEEK (6) - (0B)

Table 227: SEEK (6) - (0B)

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Command Code = 0Bh							
1	Reserved = 0			(MSB) LBA				
2	Logical Block Address (LSB)							
3								
4	Reserved = 0							
5	VU = 0		Reserved = 0				FLAG	LINK

The SEEK (6) command is implemented as a no-op on the SSD. No checking is performed on the LBA in the CDB.

17.39 SEEK (10) - (2B)

Table 228: SEEK (10) - (2B)

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Command Code = 2Bh							
1	Reserved = 0			Reserved = 0				0
2-5	(MSB) Logical Block Address (LSB)							
6-8	Reserved = 0							
9	VU = 0		Reserved = 0				FLAG	LINK

The SEEK (10) command is implemented as a no-op on the SSD. No checking is performed on the LBA in the CDB.

17.40 SEND DIAGNOSTIC (1D)

Table 229: SEND DIAGNOSTIC (1D)

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Command Code = 1Dh							
1	Function Code			PF	RSVD =0	SlfTst	Dev0fl	Unt0fl
2	Reserved = 0							
3-4	(MSB) Parameter List Length							(LSB)
5	VU = 0		Reserved = 0				FLAG	LINK

The SEND DIAGNOSTIC command requests the drive to perform its self-diagnostic test or to perform a function based on a page of information sent in a Data Out phase during the command.

- **PF (Page Format)** bit set to one indicates the data sent by the Initiator conform to the page structure as specified in SCSI standard. This bit is ignored by the Target if the SlfTst bit is set.
- **SlfTst** set to one indicates that the device performs its default self-test. If SlfTst is one, the Function code field is ignored. If SlfTst is set to zero, the action to perform is specified in Function code field.

Table 230: SEND DIAGNOSTIC Function Code (1D)

Value	Function name	Description
000b	NA	Value to be used when the SlfTst bit is set to one or if the SEND DIAGNOSTIC command is not invoking one of the other self-test function codes.
001b	Background Short self-test	The device server starts its short self-test routine in background mode.
010b	Background extended self-test	The device server starts its extended self-test routine in background mode.
011b	NA	Reserved.
100b	Abort background self-test	Abort the current self-test in the background mode. This value is only valid if a previous SEND DIAGNOSTIC command specified a background self-test function and that function has not been completed.
101b	Foreground short self-test	The device server starts its short self-test routine in the foreground mode. This self-test will complete in two minutes or less.
110b	Foreground extended self-test	The device server starts its extended self-test routine in the foreground mode. The completion time for this test is reported in Mode Page 0Ah (refer to section 17.11.9 "Mode Page 0A").
111b		Reserved.

- **DevOfI** is ignored by the Target for compatibility.
- **UntOfI** is ignored by the Target for compatibility.
- **Parameter List Length** must be 0 when the SlfTst bit is one. Otherwise, **Check Condition** status will be generated with a sense key of *Illegal Request* and additional sense of *Invalid Field in CDB*. If the SlfTst bit is zero, it should be set to the length of the page to be transferred in the DATA OUT phase of the command. If it does not match the expected length of the page a **Check Condition** status will be also generated with a sense key of *Illegal Request* and additional sense of *Invalid Field in CDB*.

If a fault is detected during the default or foreground self-test, a **Check Condition** is reported as an end status. If a fault is detected during the background self-test, it is logged in the log page for later retrieval by a LOG SENSE command.

See Section 21.19 “Diagnostics” on page 292 for a detailed listing of operations carried out by the SEND DIAGNOSTIC command and Power on Diagnostics.

17.40.1 Send Diagnostic Page 0

This page requests that the drive return a list of supported pages on the next RECEIVE DIAGNOSTICS command.

Table 231: Diagnostic Page 0

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Page Code = 0							
1	Reserved = 0							
2 - 3	Page Length = 0							

17.40.2 Send Diagnostic Page 40

This allows the Initiator to translate a LBA or physical sector address to the other format. The address to be translated is passed to the Target with the SEND DIAGNOSTIC command and the results are returned to the Initiator by the RECEIVE DIAGNOSTICS command.

The Target will read the parameter list from the Initiator, and, if no errors are detected in the parameter list, **Good** status will be returned. The data translation will be performed upon receipt of the RECEIVE DIAGNOSTICS command.

Table 232: Diagnostic Page 40

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Page Code = 40h							
1	Reserved = 0							
2-3	Page Length = 0Ah							
4	Reserved = 0				Supplied format			
5	Reserved = 0				Translate format			
6-13	Address to Translate							

Supplied Format may take one of the following three values:

- **000b** Block format
- **100b** Bytes From Index format
- **101b** Physical Sector format

It specifies the format in which the address has been supplied.

- **Translate Format** specifies the format that the address should be translated into. If the supplied format is the Block format, the Translate format must be either Bytes from Index or Physical Sector format. If the supplied format is the Bytes from Index or Physical Sector format, the Translate format must be Block format. If either of the format fields is invalid or they specify the same format, the command will terminate with **Check Condition** status with a sense code of *Illegal Request* and *Illegal Field in Parameter List*.
- **Address to Translate** contains the address to translate. If the logical block format is specified, the first four bytes of the field (bytes 6 to 9) contain the LBA and the remainder must be zero. For the physical format the address must be specified as follows.

Table 233: Address to translate

Byte	Bit							
	7	6	5	4	3	2	1	0
6-13	Physical Address							

17.40.3 Send Diagnostic Page A0h

The Device LED Control Page A0h allows the initiator to turn on or off the device fault LED light.

Table 234: Device LED Control Page - Send Diagnostic

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Page Code = A0h							
1	Reserved = 0							
2-3	(MSB) Page Length = 0002h							(LSB)
4	Reserved = 0							
5	Reserved = 0							SBDL

The page begins with a four-byte page header which specifies the page code and length. A SBDL (Set Bad Device Light) bit of one turns the device LED on, zero turns the device LED off. If the target detects an error in the passed parameter data it shall terminate the command with CHECK CONDITIONS Status. The sense key shall be set to Illegal Request and the Additional Sense code set to Illegal Field in CDB.

17.41 SET DEVICE IDENTIFIER (A4/06)

Table 235: SET DEVICE IDENTIFIER (A4/06)

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Command Code = A4h							
1	Reserved = 0				Service Action = 06h			
2	Reserved = 0							
3	Reserved = 0							
4-5	Restricted = 0							
6-9	(MSB) Parameter List Length (LSB)							
10	Reserved = 0							
11	VU = 0		Reserved = 0				FLAG	LINK

The SET DEVICE IDENTIFIER command requests that the device identifier information be set to the value received in the SET DEVICE IDENTIFIER parameter list.

On successful completion of a SET DEVICE IDENTIFIER command a unit attention is generated for all Initiators except the one that issued the service action. When reporting the unit attention condition the additional sense code is set to *Device Identifier Changed*.

- **Parameter List Length** field specifies the length in bytes of the Identifier that is transferred from the host system to the Target. The maximum value for this field is 512 bytes. A parameter list length of zero indicates that no data is transferred, and that subsequent REPORT DEVICE IDENTIFIER commands return an Identifier length of zero.

The SET DEVICE IDENTIFIER parameter list contains the identifier to be set by the addressed logical unit.

Table 236: SET DEVICE IDENTIFIER, Parameter List

Byte	Bit							
	7	6	5	4	3	2	1	0
0-n	Identifier							

The IDENTIFIER field is a vendor specific value, to be returned in subsequent REPORT DEVICE IDENTIFIER commands.

17.42 START STOP UNIT (1B)

Table 237: START STOP UNIT (1B)

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Command Code = 1Bh							
1	Reserved = 0							Immed
2	Reserved = 0							
3	Reserved = 0				Power Condition Modifier			
4	Power Condition				Reserved=0		LoEj = 0	Start
5	VU = 0		Reserved = 0				FLAG	LINK

The START STOP UNIT command is used to make the media accessible or inaccessible. At power-on, the SSD automatically makes the media accessible, but START STOP UNIT emulates HDD behaviors by making the drive ready or not ready.

- **Immed** bit is to specify
 - 0** status is to be returned at the end of the operation.
 - 1** **Good** status shall always be returned immediately after command has been received. The TEST UNIT READY command may be used to determine when the drive becomes ready.
- **Power Conditions and Power Condition Modifier fields** are ignored. Power save modes are not supported.
- **Start bit is to specify:**
 - 0** make the media inaccessible
 - 1** make the media accessible

START STOP UNIT with Start = 0 causes the firmware to save critical data to the media and make the SSD safe for power removal. Removing power without issuing START STOP UNIT with Start = 0 causes a longer time to ready on the subsequent power on.

Note: NOTIFY(ENABLE_SPINUP) is not required for the SSD to come ready after power on. Once the drive has become ready, the START STOP UNIT command can be used without any errors regardless of the current state. Note that NOTIFY(ENABLE_SPINUP) is required to transition from the Stopped state to the Ready state.

17.43 SYNCHRONIZE CACHE (10) - (35)

Table 238: SYNCHRONIZE CACHE (10) - (35)

Byte	BIT							
	7	6	5	4	3	2	1	0
0	Command Code = 35h							
1	Reserved = 0			Reserved = 0			Immed = 0	Obsolete
2-5	(MSB) Logical Block Address (LSB)							
6	Reserved = 0							
7-8	(MSB) Number of Blocks (LSB)							
9	VU = 0	Reserved = 0				FLAG	LINK	

The SYNCHRONIZE CACHE command ensures that logical blocks in the cache have their most recent data value recorded on the media.

- **Logical Block Address** is to specify where the operation is to begin.
- **Immed** (immediate) must be zero. An immediate bit of zero indicates that the status shall not be returned until the operation has completed. If the Immed bit is set to one, the drive returns a **Check Condition** status. The sense key shall be set to *Illegal Request* and the additional sense code shall be set to *Invalid Field in CDB*.
- **Number of Blocks** specifies the total number of contiguous logical blocks within the range. Number of Blocks of zero indicates that all remaining logical blocks on the logical unit shall be within the range.

17.44 SYNCHRONIZE CACHE (16) - (91)

Table 239: Synchronize Cache (16) - (91)

Byte	BIT							
	7	6	5	4	3	2	1	0
0	Command Code = 91h							
1	Reserved = 0						Immed = 0	Rsvd = 0
2-9	Logical Block Address							
10-13	Number of Blocks							
14	Reserved = 0							
15	VU = 0		Reserved = 0				FLAG	LINK

The SYNCHRONIZE CACHE command ensures that logical blocks in the cache have their most recent data value recorded on the media. See the SYNCHRONIZE CACHE (10) description for definitions of the fields in this command.

17.45 TEST UNIT READY (00)

Table 240: TEST UNIT READY (00)

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Command Code = 00h							
1	Reserved = 0			Reserved = 0				
2-4	Reserved = 0							
5	VU = 0		Reserved = 0			FLAG		LINK

The TEST UNIT READY command allows the Initiator to check if the drive is READY. The SCSI specification defines READY as the condition where the device will accept a media-access command without returning **Check Condition** status.

The drive will verify that the media is accessible. If the media is not accessible, **Check Condition** status is returned with sense key of Not Ready. If the media is accessible, the drive returns good status and will execute media access commands.

The TEST UNIT READY command is not intended as a diagnostic. No self diagnostic is performed by the device as a result of this command.

The TEST UNIT READY command has special significance for power sequencing using the UNIT START command with an Immediate bit of one. In this mode the UNIT START command returns **Task Complete** status immediately and expects the Initiator to issue TEST UNIT READY commands to determine when the media is accessible.

Note: The Power On sequence automatically transitions the drive to the Ready state. The drive does not execute any commands other than TEST UNIT READY, INQUIRY, or REQUEST SENSE command until the Power On sequence is complete. The drive will return **Check Condition** status with Not Ready sense key and In Process of Becoming Ready sense code for all other commands during the Power On period.

17.46 UNMAP (42h)

Table 241: UNMAP (42h)

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Command Code = 42h							
1	Reserved = 0							ANCHOR
2-5	Reserved = 0							
6	Reserved = 0			GROUP NUMBER				
7-8	Parameter List Length							
5	Control							

The OPERATION CODE field is defined in SPC-4 and shall be set to the value defined in Table 244:

An ANCHOR bit set to zero specifies that the LBAs on which unmap operations are performed, if any, shall become deallocated. An ANCHOR bit set to one specifies that the LBAs on which unmap operations are performed, if any, shall become anchored.

If the ANCHOR bit is set to one, and the ANC_SUP field in the Thin Provisioning VPD page is set to 000b, 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 for the definition of the GROUP NUMBER field.

The PARAMETER LIST LENGTH field specifies the length in bytes of the UNMAP parameter data that shall be sent from the application client to the device server. A PARAMETER LIST LENGTH set to zero specifies that no data shall be sent.

The contents of the CONTROL byte are defined in SAM-4.

17.46.1 UNMAP parameter list

The UNMAP parameter list contains the data sent by an application client along with an UNMAP command. Included in the data are an UNMAP parameter list header and block descriptors for LBA extents to be processed by the device server for the UNMAP command. The LBAs specified in the block descriptors may contain overlapping extents, and may be in any order.

If the ANCHOR bit in the CDB is set to zero, then each specified LBA should become deallocated and may become anchored.

If the ANCHOR bit in the CDB is set to one, and the ANC_SUP field in the Thin Provisioning VPD page is set to 001b, then for each specified LBA:

- a. If the LBA is mapped, then that LBA should become anchored, and the LBA shall not become deallocated;
- b. If the LBA is deallocated, then that LBA shall become anchored. If a lack of LBA mapping resources prevents the LBA from becoming anchored, then the command shall be terminated or
- c. If the LBA is anchored, then that LBA shall remain anchored.

Table 242: UNMAP parameter list

Byte	Bit							
	7	6	5	4	3	2	1	0
0-1	UNMAP DATA LENGTH (n-1)							
2-3	UNMAP BLOCK DESCRIPTOR DATA LENGTH (n-7)							
4-7	Reserved = 0							
8-23	UNMAP block descriptor (first)							
...							
(n-15) -n	UNMAP block descriptor (last)							

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 data. This condition shall not be considered an error.

Table 243: UNMAP block descriptor

Byte	Bit							
	7	6	5	4	3	2	1	0
0-7	UNMAP LOGICAL BLOCK ADDRESS							
8-11	NUMBER OF LOGICAL BLOCKS							
12-15	Reserved = 0							

The UNMAP LOGICAL BLOCK ADDRESS field contains the first LBA of the UNMAP block descriptor to be unmapped. The NUMBER OF LOGICAL BLOCKS field contains the number of LBAs 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 LBA specified by the UNMAP LOGICAL BLOCK ADDRESS field plus the number of logical blocks 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.

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, 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

17.47 VERIFY (2F)

Table 244: VERIFY (2F)

Byte	BIT							
	7	6	5	4	3	2	1	0
0	Command Code = 2Fh							
1	VRPROTECT		DPO	Reserved = 0			Byte Chk	RSVD = 0
2-5	(MSB) Logical Block Address (LSB)							
6	Reserved = 0							
7-8	(MSB) Verification Length (LSB)							
9	VU = 0		Reserved = 0			FLAG	LINK	

The VERIFY command requests that the drive verify the data written on the media. A verification length of zero indicates that no data will be transferred. This condition is not considered an error.

- **ByteChk** bit set to zero indicates that the data is read from the drive and verified using ECC. If an ECC error is detected in the verify process, **Check Condition** status is returned with sense key set to *Medium Error*. ByteChk bit set to one indicates that byte-by-byte comparison is performed between the data on the drive and data transferred from the initiator during the data-out phase.
If the comparison is unsuccessful, the command is terminated with **Check Condition** status and the sense key is set to *Miscompare*.
- **DPO** (Disable Page Out) bit is ignored.
- The command implies FUA.

The command stops on *Check Condition* and reports the LBA in error. The command must be reissued, starting with the next LBA, to verify the remainder of the Drive.

The Verification Length is the number of blocks to check.

The data (if any) from the data-out phase and the data from the media are not retained in the cache. Therefore, the DPO bit has no effect on this command and is ignored.

VRPROTECT defines the manner in which protection information read from drive shall be checked during processing of the command. Protection information is stored on drive, and may be validated using the drive's internal checking algorithms, and also byte-by-byte compared using data from the initiator when ByteChk=1.

If the drive is not formatted with protection information, VRPROTECT must be set to 000b, else Check Condition status will be returned with sense key of Illegal Request and additional sense code of Invalid Field in CDB.

VRPROTECT=000b

If the drive is not formatted with protection information, only user data is verified.

If the drive is formatted with protection information:

- Logical Block Guard is checked
- Logical Block Application Tag is checked (applies to VERIFY(32) command only)
- Logical Block Reference Tag is checked

VRPROTECT=001b

- Logical Block Guard is checked
- Logical Block Application Tag is checked (applies to VERIFY(32) command only)
- Logical Block Reference Tag is checked

VRPROTECT=010b

- Logical Block Guard is not checked
- Logical Block Application Tag is checked (applies to VERIFY(32) command only)
- Logical Block Reference Tag is checked

VRPROTECT=011b

- Logical Block Guard is not checked
- Logical Block Application Tag is not checked
- Logical Block Reference Tag is not checked

VRPROTECT=100b

- Logical Block Guard is checked
- Logical Block Application Tag is not checked
- Logical Block Reference Tag is not checked

VRPROTECT=101b, 110b, 111b

These values are reserved. Check Condition status will be returned with sense key of Illegal Request and additional sense code of Invalid Field in CDB.

If a check of the protection information fails, Check Condition status will be returned with sense key of Aborted Command and additional sense code indicating which protection field check failed.

If ByteChk=1, the drive's internal checking of protection information is done only when VRPROTECT=000b and the drive is formatted with protection information

If ByteChk=1, and VRPROTECT is not set to 000b, checking of protection information is performed on the fields described above as a byte-by-byte comparison against the data transferred to the drive by the initiator during the Data Out phase.

Refer to the ANSI T10 standards for additional details of protection information.

17.48 VERIFY (12) - (AF)

Table 245: Verify (12) - (AF)

Byte	BIT							
	7	6	5	4	3	2	1	0
0	Command Code = AFh							
1	VRPROTECT		DPO	FUA	Reserved=0	Byte Chk	Reserved = 0	
2-5	(MSB) Logical Block Address (LSB)							
6 - 9	(MSB) Verification Length (LSB)							
10	Reserved = 0							
11	VU = 0		Reserved = 0			FLAG	LINK	

The VERIFY(12) command causes the drive to verify data written on the media. See the VERIFY(10) description for the definitions of the fields in this command.

17.49 VERIFY (16) - (8F)

Table 246: Verify (16) - (8F)

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Command Code = 08Fh							
1	VRPROTECT		DPO	Reserved = 0			Byte Chk	Rsvd = 0
2-9	(MSB) Logical Block Address							(LSB)
10-13	(MSB) Verification Length							(LSB)
14	Reserved = 0							
15	VU = 0		Reserved = 0			FLAG	LINK	

The VERIFY command requests that the drive verify the data written on the media. See the VERIFY (10) description for the definitions of the fields in this command.

17.50 VERIFY (32) - (7F/0A)

Table 247: Verify (32) - 7F/0A)

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Command Code = 07Fh							
1	VU = 0		Reserved = 0				FLAG	LINK
2-5	Reserved = 0							
6	Reserved = 0			Group Number = 0				
7	Additional CDB Length = 18h							
8 - 9	Service Action = 000Ah							
10	RDPROTECT			DPO	Reserved = 0		ByteC hk	Reserv ed=0
11	Reserved = 0							
12 - 19	(MSB) Logical Block Address (LSB)							
20 - 23	(MSB) Expected Initial Logical Block Reference Tag (LSB)							
24 - 25	(MSB) Expected Logical Block Application Tag (LSB)							
26 - 27	(MSB) Logical Block Application Tag Mask (LSB)							
28 - 31	(MSB) Verification Length (LSB)							

The VERIFY command requests that the verify the data written on the media. Each logical block includes user data and may include protection information, based on the VPROTECT field and the drive format.

If the drive is formatted with type 2 protection (PROT_EN=1 and P_TYPE=001b in the READ CAPACITY (16) parameter data), then this command will be processed normally. Any other protection types will result in Check Condition status to be returned with sense key of Illegal Request and additional sense code of Invalid Command Operation Code

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.

If the ATO bit is set to one in Mode Page 0Ah, 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 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 the protection information.

If the ATO bit is set to zero, the Logical Block Application Tag Mask field and the Expected Logical Block Application Tag field are ignored.

17.51 WRITE (6) - (0A)

Table 248: WRITE (6) - (0A)

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Command Code = 0Ah							
1	Reserved = 0			(MSB) LBA				
2-3	Logical Block Address (LSB)							
4	Transfer Length							
5	VU = 0		Reserved = 0				FLAG	LINK

The WRITE command requests the drive to write the specified number of blocks of data (**Transfer Length**) from the Initiator to the medium starting at the specified **Logical Block Address (LBA)**.

See Section 17.15 READ (6) - (08)” on page 188 for the parameters.

17.52 WRITE (10) - (2A)

Table 249: WRITE (10) - (2A)

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Command Code = 2Ah							
1	WRPROTECT			DPO	FUA	Rsvd=0	FUA_N V	Obso- lete
2-5	(MSB) Logical Block Address (LSB)							
6	Reserved = 0							
7-8	(MSB) Transfer Length (LSB)							
9	VU = 0	Reserved = 0				FLAG	LINK	

The WRITE (10) command requests that the drive write the data transferred from the Initiator. This command is processed like the standard WRITE (6) - (0A) command except for the longer transfer length.

- **Transfer Length** is the number of contiguous blocks to be transferred. If the transfer length is zero, the seek occurs, but no data is transferred. This condition is not considered an error.
- **DPO** (Disable Page Out) bit is ignored.
- **FUA** (bit is ignored).
- **FUA_NV** (Force Unit Access Non-Volatile Cache) may be set to 0 or 1, but is ignored since NV_SUP=0 in Inquiry Page 86h.

If a WRITE(6) command is received after protection information is enabled, the drive will set the protection information as follows as it writes each block to drive:

- the Logical Block Guard field is set to a properly generated CRC
- the Logical Block Reference Tag field is set to:
 - the least significant four bytes of the LBA, if the drive is formatted with type 1 protection (PROT_EN=1 and P_TYPE=000b in the READ CAPACITY (16) parameter data); or
 - FFFFFFFFh, if the drive is formatted with type 2 protection (PROT_EN=1 and P_TYPE=001b in the READ CAPACITY (16) parameter data)
- the Logical Block Application Tag field is set to
 - FFFFh, if the ATO bit is set to one in Mode Page 0Ah; or
 - Any value, if the ATO bit is set to zero

WRPROTECT defines the manner in which protection information written to drive shall be checked during processing of the command. Protection information may be transmitted to the drive with the user data, based on the WRPROTECT bit and the drive format.

If the drive is not formatted with protection information, WRPROTECT must be set to 000b, else Check Condition status will be returned with sense key of Illegal Request and additional sense code of Invalid Field in CDB.

WRPROTECT=000b

Protection information is not transmitted to the drive.

If the drive is formatted with protection information, the drive will write protection information to drive based on its internal algorithms.

WRPROTECT=001b

- Protection information is transmitted to the drive with the user data
- Logical Block Guard is checked
- Logical Block Application Tag is checked (applies to WRITE (32) command only)
- Logical Block Reference Tag is checked

WRPROTECT=010b

- Protection information is transmitted to the drive with the user data
- Logical Block Guard is not checked
- Logical Block Application Tag is checked (applies to WRITE(32) command only)
- Logical Block Reference Tag is checked

WRPROTECT=011b

- Protection information is transmitted to the drive with the user data
- Logical Block Guard is not checked
- Logical Block Application Tag is not checked
- Logical Block Reference Tag is not checked

WRPROTECT=100b

- Protection information is transmitted to the drive with the user data
- Logical Block Guard is checked
- Logical Block Application Tag is not checked
- Logical Block Reference Tag is not checked

WRPROTECT=101b, 110b, 111b

These values are reserved. Check Condition status will be returned with sense key of Illegal Request and additional sense code of Invalid Field in CDB.

If a check of the protection information fails, Check Condition status will be returned with sense key of Aborted Command and additional sense code indicating which protection field check failed.

Refer to the ANSI T10 standards for additional details of protection information.

17.53 WRITE (12) - (AA)

Table 250: Write (12) - (AA)

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Command Code = AAh							
1	WRPROTECT			DPO	FUA	Rsvd=0	FUA_NV	Rsvd=0
2-5	(MSB) Logical Block Address (LSB)							
6-9	(MSB) Transfer Length (LSB)							
10	Reserved=0							
11	VU = 0		Reserved = 0				FLAG	LINK

The WRITE(12) command causes the drive to write data from the initiator to the media. See the WRITE(10) description for the definitions of the fields in this command.

17.54 WRITE (16) - (8A)

Table 251: Write (16) - (8A)

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Command Code = 8Ah							
1	WRPROTECT			DPO	FUA	Rsvd=0	FUA_NV	Rsvd=0
2-9	(MSB) Logical Block Address (LSB)							
10-13	(MSB) Transfer Length (LSB)							
14	Reserved = 0							
15	VU = 0		Reserved = 0				FLAG	LINK

The WRITE(16) command causes the drive to write data from the initiator to the media. See the WRITE(10) description for the definitions of the fields in this command.

17.55 WRITE (32) - (7F/0B)

Table 252: Write (32) - (7F/0B)

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Command Code = 7Fh							
1	VU = 0		Reserved = 0				FLAG	LINK
2-5	Reserved = 0							
6	Reserved = 0			Group Number = 0				
7	Additional CDB Length = 18h							
8-9	Service Action = 000Bh							
10	WRPROTECT			DPO	FUA	Rsvd=0	FUA_NV	Rsvd=0
11	Reserved = 0							
12-19	(MSB)	Logical Block Address						(LSB)
20-23	(MSB)	Expected Initial Logical Block Reference Tag						(LSB)
24-25	(MSB)	Expected Logical Block Application Tag						(LSB)
26-27	(MSB)	Logical Block Application Tag Mask						(LSB)
28-31	(MSB)	Transfer Length						(LSB)

The WRITE command requests that the drive write data transferred from the initiator to drive. Each logical block transferred

includes user data and may include protection information, based on the WRPROTECT field and the drive format. Each logical block written includes user data and, if the drive is formatted with protection information enabled, protection information. If the drive is formatted with type 2 protection (PROT_EN=1 and P_TYPE=001b in the READ CAPACITY (16) parameter data), then this command will be processed normally. Any other protection types will result in Check Condition status to be returned with sense key of Illegal Request and additional sense code of Invalid Command Operation Code

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.

If the ATO bit is set to one in Mode Page 0Ah, 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 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 the protection information.

If the ATO bit is set to zero, the Logical Block Application Tag Mask field and the Expected Logical Block Application Tag field are ignored.

17.56 WRITE AND VERIFY (10) - (2E)

Table 253: WRITE AND VERIFY (10) - (2E)

Byte	Bit								
	7	6	5	4	3	2	1	0	
0	Command Code = 2Eh								
1	WRPROTECT			DPO	Reserved = 0			Byte Chk	Obsolete
2-5	(MSB) Logical Block Address (LSB)								
6	Reserved = 0								
7-8	(MSB) Transfer Length (LSB)								
9	VU = 0		Reserved = 0				FLAG	LINK	

WRITE AND VERIFY command requests that the drive writes the data transferred from the Initiator to the medium and then verify that the data is correctly written. An implied FUA (Force Unit Access) and an implied Synchronize Cache are performed before starting the operation. This insures that data from the drive, not the cache, is verified.

- See the WRITE (10) command description for the definition of the WRPROTECT field.
- **Transfer Length** is the number of contiguous blocks to transferred. If the transfer length is zero, the seek occurs, but no data is transferred. This condition is not considered an error.
- **ByteChk** bit set to zero indicates that the data is read back from the drive and verified using ECC after the successful write operation. If an ECC error is detected in the verify process, **Check Condition** status is returned with sense key set to *Medium Error*. ByteChk bit set to one indicates that byte-by-byte comparison is performed between data on the drive starting the block specified in LBA field and data transferred from the Initiator.

If the comparison is unsuccessful, the command is terminated with **Check Condition** status and the sense key is set to *Miscompare*.

- **DPO** (Disable Page Out) bit is ignored.

17.57 WRITE AND VERIFY (12) - (AE)

Table 254: Write and Verify (12) - (AE)

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Command Code = AEh							
1	WRPROTECT			DPO	Reserved = 0		ByteChk	Obsolete
2-5	(MSB) Logical Block Address (LSB)							
6-9	(MSB) Transfer Length (LSB)							
10	Reserved = 0							
11	VU = 0		Reserved = 0			FLAG	LINK	

The WRITE AND VERIFY command requests that the drive write the data transferred from the Initiator to the medium and then verify that the data is correctly written. See the WRITE AND VERIFY (10) description for the definitions of the fields in this command.

17.58 WRITE AND VERIFY (16) - (8E)

Table 255: Write and Verify (16) - (8E)

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Command Code = 8Eh							
1	WRPROTECT			DPO	Reserved = 0		Byte Chk	Obsolete
2-9	(MSB) Logical Block Address (LSB)							
10-13	(MSB) Transfer Length (LSB)							
14	Reserved = 0							
15	VU = 0		Reserved = 0			FLAG	LINK	

The WRITE AND VERIFY command requests that the drive write the data transferred from the Initiator to the medium and then verify that the data is correctly written.

17.59 WRITE AND VERIFY (32) - (7F/0C)

Table 256: Write and Verify (32) - (7F/0C)

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Command Code = 7Fh							
1	VU = 0		Reserved = 0				FLAG	LINK
2 - 5	Reserved = 0							
6	Reserved = 0			Group Number = 0				
7	Additional CDB Length = 18h							
8 - 9	Service Action = 000Ch							
10	WRPROTECT			DPO	Reserved = 0		ByteChk	Reserved=0
11	Reserved = 0							
12 - 19	Logical Block Address							
20 - 23	Expected Initial Logical Block Reference Tag							
24 - 25	Expected Logical Block Application							
26 - 27	Logical Block Application Tag Mask							
28 - 31	Transfer Length							

The WRITE AND VERIFY command requests that the drive write the data transferred from the initiator to drive and then verify that the data is correctly written.

If the drive is formatted with type 2 protection (PROT_EN=1 and P_TYPE=001b in the READ CAPACITY (16) parameter data), then this command will be processed normally. Any other protection types will result in Check Condition status to be returned with sense key of Illegal Request and additional sense code of Invalid Command Operation Code

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.

If the ATO bit is set to one in Mode Page 0Ah, 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 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 the protection information.

If the ATO bit is set to zero, the Logical Block Application Tag Mask field and the Expected Logical Block Application Tag field are ignored.

17.60 WRITE BUFFER (3B)

Table 257: WRITE BUFFER (3B)

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Command Code = 3Bh							
1	Reserved = 0			Mode				
2	Buffer ID							
3-5	Buffer Offset							
6-8	Parameter List Length							
9	VU = 0		Reserved = 0				FLAG	LINK

The WRITE BUFFER command is used in conjunction with the READ BUFFER command as a diagnostic function for testing the memory of the drive and the SCSI bus integrity. This command does not alter the medium of the drive. Additional modes are provided for downloading microcode and saving microcode.

This command will cause the entire cache to be emptied.

The function of this command and the meaning of fields within the command descriptor block depend on the contents of the mode field.

MODE	Description
00000	Write combined header and data
00010	Data
00100	Download Microcode
00101	Download Microcode and Save - single binary file
00111	Download Microcode and Save - multiple binary files
01010	Write Data to Echo Buffer
11010	Enable expander Communications Protocol
All Others	Not Supported

If any values other than shown above are specified, **Check Condition** status is returned with a sense key of *Illegal Request* and additional sense code of *Invalid Field in CDB*.

17.60.1 Combined Header And Data (Mode 00000b)

In this mode, the data to be transferred is preceded by a four-byte header.

Buffer ID must be zero. If another value is specified, no download function is performed and the command is terminated with **Check Condition** status. And the drive shall set the sense key to *Illegal Request* and additional sense code to *Illegal Field in CDB*.

Buffer Offset must be zero. If another value is specified, no download function is performed and the command is terminated with **Check Condition** status. And the drive shall set the sense key to *Illegal Request* and additional sense code to *Illegal Field in CDB*.

Parameter List Length specifies the number of bytes that shall be transferred during the DATA OUT phase. This number includes four bytes of header, so the data length to be stored in the drive buffer is transfer length minus four. If the length

exceeds the buffer size, the command is terminated with **Check Condition** status. And the drive shall set sense key to *Illegal Request* and additional sense code to *Illegal Field in CDB*. A Parameter List Length of less than four (size of header) indicates no data is transferred.

The 4-byte header consists of all reserved bytes.

Table 258: Write Buffer Header

Byte	Bit							
	7	6	5	4	3	2	1	0
0-3	Reserved = 0							

17.60.2 Write Data (Mode 00010b)

In this mode, the DATA OUT phase contains buffer data.

Buffer ID must be zero. If another value is specified, no download function is performed and the command is terminated with Check Condition status. And the drive shall set the sense key to *Illegal Request* and additional sense code to *Illegal Field In CDB*.

Buffer Offset specifies the offset of the memory space specified by the Buffer ID. The initiator should conform to the offset boundary requirements returned in the READ BUFFER descriptor. If the value exceeds the buffer specified, the command is terminated with **Check Condition** status. And the drive shall set the sense key to *Illegal Request* and additional sense code to *Illegal Field In CDB*.

Parameter List Length specifies the Parameter List Length. It must be

- less than the capacity of the buffer size after adding the Buffer Offset value and
- on a sector boundary

A Parameter List Length of zero indicates no data is to be transferred and command status is returned.

If an invalid value is specified, the command is terminated with **Check Condition** status. And the drive shall set the sense key to *Illegal Request* and additional sense code to *Illegal Field In CDB*.

17.60.3 Download Microcode (Mode 00100b)

NOTE: It is not expected that a customer will ever issue this format of the command.

In this mode, the microcode is transferred to the control memory space of the drive. When downloaded, the drive will operate with the newly downloaded code immediately until the next power cycle.

Buffer ID field is used to indicate which portion of the microcode image is being downloaded. The following Buffer IDs are supported by the Target:

- 00h: Main Microprocessor Code
- nnh: ID of Vendor Unique Reserved Area

Any unsupported value for the Buffer ID will cause the command to terminate with **Check Condition** status. And the drive shall set the sense key to *Illegal Request* and additional sense code to *Illegal Field In CDB*.

Buffer Offset must be zero. If an invalid value is specified, the command is terminated with **Check Condition** status. The drive shall set the sense key to *Illegal Request* and additional sense code to *Illegal Field in CDB*.

Parameter List Length must be the size of the data set to be downloaded. It may also be set to 0000h in which case no code is updated and command status is returned. If an invalid value is specified, the command is terminated with **Check Condition** status. And the drive shall set the sense key to *Illegal Request* and additional sense code to *Illegal Field In CDB*.

This process generates a unit attention condition for MICROCODE HAS BEEN CHANGED for all Initiators except the one which sent the WRITE BUFFER command. Upon the completion of the WRITE BUFFER command the new microcode is

immediately ready for operation.

Note: The Download Microcode mode described in this specification is to indicate that the drive will accept a command with this mode, though it is not expected that a user will ever issue such a command. To use the write buffer command with this mode, a special microcode version is required from development. If such a microcode is released from development, then it will include appropriate instructions on the function of new microcode and its effect on the drive operations after download.

17.60.4 Download Microcode and Save (Mode 00101b) -Single Binary File

In this mode the data is transferred to the drive to save into the System reserved area on the drive. This is for functional upgrade and configuration change reflecting the user's requirements and the manufacturer's reason or both, and it is stored in the media as a permanent copy. The newly downloaded code becomes effective after the drive issues and completes a self-initiated Power On Reset.

Note: It requires up to 30 seconds to update the microcode including the Flash ROM update.

Note: New code to be downloaded to the drive will be provided by development either by request of a customer for an additional function or as a result of a functional change by development. However please note that not all possible fixes or new functions can be applied to a drive in this manner and that there is a very high dependency on the level of ROM code contained within the drive. If an invalid code or a code not compatible with the ROM code is downloaded, the drive will usually reject this code and will continue normal operation. However there is a small possibility that an invalid code will be accepted. If this occurs, the unit usually becomes inoperable and will have to be returned to the manufacturer for recovery.

Buffer ID field is used to indicate which portion of the microcode image is being downloaded. To download microcode, the buffer ID should be set to 00h. Other values are reserved for HGST development purposes only.

17.60.5 Download Microcode and Save (Mode 00111b) - Multiple Binary Files

In this mode the target receives a segment of the binary microcode file. The Parameter List Length (segment length) of each segment shall be a multiple of 4K bytes. The total length of all segments received shall be equal to the total length of the binary microcode file. All segments must be sent in the proper sequential order.

If an invalid Parameter List Length is specified, **Check Condition** status is returned with sense key of Illegal Request and additional sense code of *Invalid Field in CDB*.

The first segment sent in this mode indicates, by default, the first segment of the binary microcode file. If a **Check Condition** status is returned in this mode, a **Buffer ID** == 00h in the subsequent Write Buffer command in this mode indicates the first segment of the binary microcode file. Otherwise the **Buffer ID** field is ignored.

The **Buffer Offset** field is ignored.

After all segments of the binary microcode file have been received, the drive behavior is the same as Download Microcode and Save (Mode 00101b) - Single Binary File.

17.60.6 Write Data to Echo Buffer (Mode 01010b)

In this mode the Target transfers data into the echo buffer. The echo buffer is assigned in the same manner by the Target as it would for a WRITE operation. Data will be sent aligned on 4-byte boundaries.

Upon successful completion of a WRITE BUFFER command the data will be preserved in the echo buffer unless there is an intervening command to any logical unit, in which case it may be changed.

17.60.7 Enable Expander Communications Protocol (Mode 11010b)

In this mode the drive behavior is the same as Write Data to Echo Buffer (Mode 01010b).

17.61 WRITE LONG (3F)

Table 259: WRITE LONG (3F)

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Command Code = 3Fh							
1	COR_DIS	WR_UNCOR	PBLOCK =0	Reserved = 0				Obsolete
2-5	(MSB) Logical Block Address							(LSB)
6	Reserved = 0							
7-8	(MSB) Byte Transfer Length							(LSB)
9	VU = 0		Reserved = 0				FLAG	LINK

The WRITE LONG command requests the drive to write **one block** of data transferred from the Initiator.

The transfer data must include

- User Data
- 4 bytes of CRC data

Parameters are

- **COR_DIS** - correction disabled, bit 7 in byte 1. When this bit is set to one, we mark the LBA as a pseudo unrecovered error with correction disabled. A subsequent read to this LBA would:
 - a) Perform no error recovery on the block;
 - b) Perform no automatic reallocation of the affected logical blocks, including any automatic reallocation enabled by the Read-Write Error Recovery mode page;
 - c) Not consider errors on the affected logical blocks to be informational exception conditions as defined in the Information Exceptions Control mode page (see SPC-4);
 - d) not log errors on the affected logical blocks in the Error Counter log pages
 - e) On a read to the LBA, return check condition status with the sense key set to Medium Error and the additional sense code set to read error marked bad by client.
- **WR_UNCOR** - write uncorrectable, bit 6 in byte 1. If we receive a Write Long command with the WR_UNCOR bit set to one, we would create a pseudo unrecovered error with correction enabled. On following read commands to the LBA, the drive will:
 - a) use our normal recovery procedures (which will end in a hard error);
 - b) perform no automatic reallocation of the affected logical blocks, including any automatic reallocation enabled by 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 SPC-4);

- d) log errors on the affected logical blocks in the Error Counter log pages
- e) On a read to the LBA, return check condition status with the sense key set to Medium Error and the additional sense code set to read error marked bad by client.

The error state for LBA written with the COR_DIS or WR_UNCOR bits set, will remain in effect until the LBA is rewritten by a write, write same, format, write long without COR_DIS set, reassign or write verify command.

- **Logical Block Address** field specifies the logical block at which the write operation shall occur.
- **Byte Transfer Length.** This field must specify the exact number of bytes of data available for transfer. If a non-zero byte transfer length does not match the available data length, the Target terminates the command with **Check Condition** status, then the sense key is set to *Illegal Request*, and an additional sense code is set to *Invalid Field in CDB*. The valid and ILI bits are set to one and the information field is set to the difference of the requested length minus the actual length in bytes. Negative values are indicated by two's complement notation.

Note: Since the ECC bytes are not included in the Read Long data, ECC correction capability cannot be tested using Read/Write Long. In addition to using COR_DIS or WR_UNCOR, a hard error can be emulated by modifying the Read Long data and writing it using Write Long. The CRC returned in the Read Long data transfer is not re-calculated on a Write Long, and a CRC check failure on a subsequent Read will return check condition status with the sense key set to Medium Error and the additional sense code set to UNRECOVERED READ ERROR. th in bytes. Negative values are indicated by two's complement notation.

17.62 WRITE SAME (41)

Table 260: WRITE SAME (41)

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Command Code = 41h							
1	WRPROTECT			Reserved = 0		PBDA TA=0	LBDA TA=0	Obso- lete
2-5	(MSB) Logical Block Address							(LSB)
6	Reserved = 0							
7-8	(MSB) Number of Blocks							(LSB)
9	VU = 0	Reserved = 0				FLAG	LINK	

The WRITE SAME command instructs the Target to write a single block of data transferred to the Target from the Initiator to a number of sequential logical blocks. This command is useful for writing large data areas without sending all of the data over the SCSI bus.

- See the WRITE(10) command description for the definition of the WRPROTECT field.
- **Logical Block Address** specifies the address at which the write begins. The Number of Blocks specifies the number of contiguous blocks to be written. If the number is zero, all of the remaining blocks on the specified Logical Unit are written.
- **Number of Blocks** specifies the number of contiguous blocks to be written. If the number is zero, all of the remaining blocks on the specified logical unit are written.
- **RelAdr** (Relative Block Address) is not supported and must be set to be zero.

The data for this command is not retained in the cache.

17.63 WRITE SAME (16) - (93)

Table 261: Write Same (16) - (93)

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Command Code = 93h							
1	WRPROTECT		ANCHOR	UNMAP	PBDATA =0	LBDATA =0	Reserved	
2-9	(MSB) Logical Block Address							(LSB)
10-13	(MSB) Number of Blocks							(LSB)
14	Reserved = 0							
15	VU = 0		Reserved = 0			FLAG	LINK	

The Write Same command instructs the Target to write a single block of data transferred to the Target from the Initiator to a number of sequential logical blocks. See the WRITE(10) command description for the definition of the WRPROTECT field.

Table 262: ANCHOR bit, UNMAP Bit, and ANC_SUP bit

UNMAP Bit	ANCHOR Bit	ANC_SUP Field	Action
0b	0b	n/a	Write
	1b	n/a	Error
1b	0b	n/a	Unmap
	1b	000b	Error
		001b	Anchor

For Action:

- **Write:** The device server shall perform the specified write operation on each LBA specified by the command.
- **Error:** 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.
- **Unmap:** The device server should deallocate each LBA specified by the command but may anchor each LBA specified by the command. If the device server does not deallocate or anchor the LBA, then the device server shall perform the specified write operation.
- **Anchor:** The device server should anchor each LBA specified by the command instead of performing the specified write operation. If the device server does not anchor the LBA, the device server shall perform the specified write operation.

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:

- d. The logical unit is fully provisioned (i.e., the TPE bit is set to zero in the READ CAPACITY (16) parameter data; and
- e. The UNMAP bit is set to one or the ANCHOR bit is set to one.

17.64 WRITE SAME (32) - (7F/0D)

Table 263: Write Same (32) - (7F/0D)

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Command Code = 7Fh							
1	VU = 0		Reserved = 0				FLAG	LINK
2-5	Reserved = 0							
6	Reserved = 0			Group Number = 0				
7	Additional CDB Length = 18h							
8 - 9	Service Action = 000Dh							
10	WRPROTECT			ANCHOR	UNMAP	PBDATA =0	LBDATA =0	Reserved
11	Reserved = 0							
12 - 19	(MSB)	Logical Block Address						(LSB)
20 - 23	(MSB)	Expected Initial Logical Block Reference Tag						(LSB)
24 - 25	(MSB)	Expected Logical Block Application Tag						(LSB)
26 - 27	(MSB)	Logical Block Application Tag Mask						(LSB)
28 - 31	(MSB)	Number of Blocks						(LSB)

The WRITE SAME command requests that the drive write a single block of data transferred from the initiator to drive for a number of sequential logical blocks. This command is useful for writing large data areas with the same data, without sending all of the data over the interface. Each logical block transferred includes user data and may include protection information, based on the WRPROTECT field and the drive format. Each logical block written includes user data and, if the drive is formatted with protection information enabled, protection information.

If the drive is formatted with type 2 protection (PROT_EN=1 and P_TYPE=001b in the READ CAPACITY (16) parameter data), then this command will be processed normally. Any other protection types will result in Check Condition status to be returned with sense key of Illegal Request and additional sense code of Invalid Command Operation Code

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.

If the ATO bit is set to one in Mode Page 0Ah, 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 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 the protection information.

If the ATO bit is set to zero, the Logical Block Application Tag Mask field and the Expected Logical Block Application Tag field are ignored.

18.0 SCSI Status Byte

Upon the completion of a command a status byte is sent to the initiator. Additional sense information may also be available depending on the contents of the status byte. The following section describes the possible values for the status byte and sense data. All Reserved fields are set to zero.

Table 264: SCSI Status Byte. Format of the SCSI STATUS byte.

Bit							
7	6	5	4	3	2	1	0
Reserved = 0		Status Code					RSVD

STATUS BYTE Description

00h	GOOD The command has been successfully completed.
02h	CHECK CONDITION An error, exception, or abnormal condition has been detected. The sense data is set by the drive. The REQUEST SENSE command should be issued to determine the nature of the condition.
04h	CONDITION MET This status is returned when an unlinked PRE-FETCH command has been successfully completed.
08h	BUSY This condition is returned when disconnect privilege is not granted while the drive is BUSY processing the other command for the other initiator. The normal initiator recovery action is to issue the command at a later time or to reissue the command and grant the disconnect privilege.
10h	INTERMEDIATE Not supported.
14h	INTERMEDIATE CONDITION MET Not supported.
18h	RESERVATION CONFLICT This status is returned whenever an SCSI device attempts to access the drive, but it has been reserved by another initiator.
28h	QUEUE FULL This status indicates that the target's command queue is full. If a tagged command queuing feature is enabled and there is no room on the command queue, this status is returned when the initiator sends a command. For this status, sense data are not valid.

19.0 Additional information

This chapter provides additional information or descriptions of various functions, features, or operating models supported by the Target that are not fully described in previous chapters.

19.1 Obtaining an AL_PA

The SCA-2 FC-AL connector used by the drive provides 7 bits which encode a preferred Hard AL_PA for use on both loops. If this value is set to 7Fh, a Soft AL_PA will be used.

Note: Loop Initialization cannot be bypassed if Soft AL_PA is selected.

A Soft AL_PA will also be used if, for any reason, the Hard AL_PA cannot be obtained. If this happens, The drive will

- log out all initiators
- accept new port logins
- accept each process login with a reason code of: 05h

No special action is taken if the connector supplied AL_PA is changed while the drive is participating on the loop(s).

19.2 Loop Initialization Procedure

L_Ports use the Loop Initialization Procedure in order to acquire an AL_PA and participate on the loop.

The drive optionally initiates the Loop Initialization Procedure shortly after power-on, depending on the configured mode settings. In addition, it may (depending on the mode settings) initiate the Loop Initialization Procedure for the following reasons:

3. When it has been unable to obtain the loop for E_D_TOV.
4. When a loop failure has been detected (loss of sync for R_T_TOV)

The Loop Initialization Procedure is disruptive to any communication that is actually in progress on the loop. IOPs that are queued or active but without any frames 'in flight' are not affected.

Initiators are not logged out by the Loop Initialization Procedure, but it is their responsibility to login if their AL_PA changes. The drive detects (via the Port_Name) this situation and implicitly ends the prior login. The initiator may perform a 'Port Discovery' sequence to check that the targets have not physically changed.

The L_Port that initiates the Loop Initialization Procedure must transmit LIPs for 2 x AL_TIME to ensure that all L_Ports participate. By default, AL_TIME is set to 15ms.

All L_Ports must provide sufficient buffering to receive and transmit the LIxx frames that are circulated during the Loop Initialization Procedure. The largest of these (LILP and LIRP) have a 132 byte payload and are supported by the drive.

Table 265: LIXx Frame Payloads

LIXx Frame	Payload	
LISM	1101 0000h	8 - byte Port_Name
LIFA	1102 0000h	16 - byte AL_PA bit map
LIPA	1103 0000h	16 - byte AL_PA bit map
LIHA	1104 0000h	16 - byte AL_PA bit map
LISA	1105 0000h	16 - byte AL_PA bit map
LIRP	1106 0000h	128 - byte AL_PA position map
LILP	1107 0000h	128 - byte AL_PA position map

The Loop Initialization Procedure is as follows...

1. Select Initial AL_PA
The initial AL_PA is EFh.

2. Select a Loop Master

LISM frames are continuously transmitted. The D_ID and S_ID fields are set to 0000EFh.

If a LISM frame is received which is the same as the one transmitted, the Port becomes the Loop Master and continues the procedure at step 3.

For all other received LISM frames...

- If the D_ID is 000000h, the frame is retransmitted.
- If the Port_Name in the payload is greater than the Port_Name, the frame is discarded.
- If the Port_Name in the payload is less than the Port_Name, the frame is retransmitted.

If an ARB(F0) is received, the Port continues the procedure at step 4.

3. Loop Master

- a. ARB(F0) is continuously transmitted until ARB(F0) is received.
- b. The Port prepares an LIFA frame with an initial AL_PA bit map of all zeros. After transmitting the LIFA frame, the Port waits to receive a LIPA frame.
- c. The AL_PA bit map from the LIFA frame is used in a LIPA frame. If the bit corresponding to the Port's previously assigned AL_PA is not set it is set now. If it is already set, the Port attempts to obtain its Hard AL_PA in step d. The Port transmits the LIPA and waits to receive it back.
- d. The AL_PA bit map from the LIPA frame is used in a LIHA frame. If the bit corresponding to the Port's hard assigned AL_PA is not set it is set now (unless a bit was set for LIPA). If it is already set, the Port attempts to obtain a Soft AL_PA in step e. The Port transmits the LIHA and waits to receive it back.
- e. The AL_PA bit map from the LIHA frame is used in a LISA frame. If a bit was not set for LIHA, the bit associated with the lowest priority AL_PA still available is set. This is the Port's "soft assigned" AL_PA. If no bits are available, the Port shall stay in non-participating mode. The Port transmits the LISA and waits to receive it back. If byte 3 of the payload in the received LISA frame is still 01h (indicating that all nodes on the loop support positional mapping) the drive continues with the next two steps. Otherwise, the next two steps are skipped.
- f. The Port creates a LIRP frame with an AL_PA position map of all FFh, except for the first two bytes which are 01xxh (xx is the Port's AL_PA). The Port transmits the LIRP frame and waits to receive it back.
- g. The AL_PA position map from the received LIRP is used to create a LILP frame. The Port transmits the LILP frame and waits to receive it back.
- h. CLS is transmitted to put all other Ports into MONITORING state. When received back, the Loop Master itself goes to MONITORING state.

4. Non Loop Master

The Port retransmits received ARB(F0) and prepares to receive the following frames, followed by CLS.

- a. When a LIFA frame is received, the Port will set the bit corresponding to the Port's "Fabric Assigned" AL_PA if the Port was logged-in to the local FL_Port prior to the LIP. (See 19.5, "Public Loop Operation" on page 276.) After transmitting the LIFA frame, The Port waits to receive a LIPA frame.
- b. When a LIPA frame is received (and a bit was not set for LIPA), the bit corresponding to any previously assigned AL_PA is checked. If not set, it is set now. If it is already set, the Port attempts to obtain its Hard AL_PA in step c. The Port transmits the LIPA.
- c. When a LIHA frame is received (and a bit was not set for LIPA), the bit corresponding to any required hard assigned AL_PA is checked. If not set, it is set now. If it is already set, the Port attempts to obtain a Soft AL_PA in step d. The Port transmits the LIHA.
- d. When a LISA frame is received (and a bit was not set for LIHA), the bit associated with the lowest priority AL_PA still available is set. This is the Port's soft assigned AL_PA. If no bits are available, the Port shall stay in non-participating mode. The Port transmits the LISA.
- e. If a LIRP frame is received, and the Port was able to claim an AL_PA in one of the LIxA frames, the left most byte is incremented by one and used as an offset at which to store the Port's AL_PA. The Port transmits the LIRP.
- f. If a LILP frame is received, the Port may store the AL_PA position map for use in error recovery. The Port transmits the LILP.
- g. When CLS is received, the Port retransmits it and goes to MONITORING state.

19.3 Flow Control

The drive provides a Class 3 service. Flow control is maintained using R_RDYs.

Each initiator is required to login with every target. This establishes the Login BB_Credit, (i.e. how many frames the initiator or target may send immediately on opening or being opened.)

Prior to login, BB_Credit is zero.

The Alternate BB_Credit Model is operated as follows...

- On opening, the Available BB_Credit in each direction is set to the agreed Login BB_Credit value.
- Each frame sent decrements the Available BB_Credit and each R_RDY received increments it.
- If the Available BB_Credit reaches zero, the node must stop sending frames until enough R_RDYs have been received to restore positive credit.
- It is permissible to close before credit is balanced (i.e. R_RDY received for each frame sent) but only when enough buffers are available to satisfy the maximum Login BB_Credit that has been extended.

Available BB_Credit may exceed Login BB_Credit but a point to note is that R_RDY is not specific to any frame type. Therefore, the amount of Available BB_Credit that the drive will extend is limited by the amount of buffer space available for FCP_CMND or FCP_DATA IUs. For example, during the execution of a 512K WRITE command, there may be buffer space for 256 2K FCP_DATA IUs and 10 FCP_CMND IUs. The amount of credit extended will be for 10 frames.

19.4 Login Requirements

In order to communicate effectively, a pair of L_Ports must exchange operational parameters. The same applies to any SCSI processes that use the L_Ports.

The exchange of parameters is called **Login** and may be either **explicit** or **implicit**.

- **Explicit Login** is the default. L_Ports use the PLOGI Extended Link Service sequence to exchange parameters and SCSI Processes use the PRLI Extended Link Service.
- **Implicit Login** is enabled via setting of mode page 19h. When enabled, PLOGI and PRLI are not required and a default set of parameters is assumed.

The response to certain frames at various stages of login is defined in Table 266.

Table 266: Response to Frames before PLOGI or PRLI

Frame Received	No PLOGI	PLOGI but no PRLI
ABTS	Discard and Send LOGO	BA_ACC or BA_RJT(1)
ADISC	Discard and Send LOGO	ACC(2) or LS_RJT
LOGO	ACC	ACC
PDISC	Discard and Send LOGO	ACC(2) or LS_RJT
PLOGI	ACC	ACC
PRLI	Discard and Send LOGO	ACC
PRLO	Discard and Send LOGO	ACC or LS_RJT(3)
Unlisted Extended Link Service	Discard and Send LOGO	ACC or LS_RJT(4)
Note: 1 - The drive returns BA_RJT to ABTS BLS when ABTS BLS has RX_ID other than 0FFFFh 2 - N_Port ID, Port_Name and Node Name must match a logged in port for ACC to be returned. 3 - If PRLI has not been successfully completed, Reason Code = 'Image Pair does not exist' 4 - Unsupported Extended Link Services are rejected.		

Login parameters remain in effect until reset or updated by...

- Power cycle
- LIP Reset
- Explicit Logout (LOGO)
- Explicit Login (PLOGI/PRLI)
- Implicit Logout (failure to re-Discover within RR_TOV following LIP)

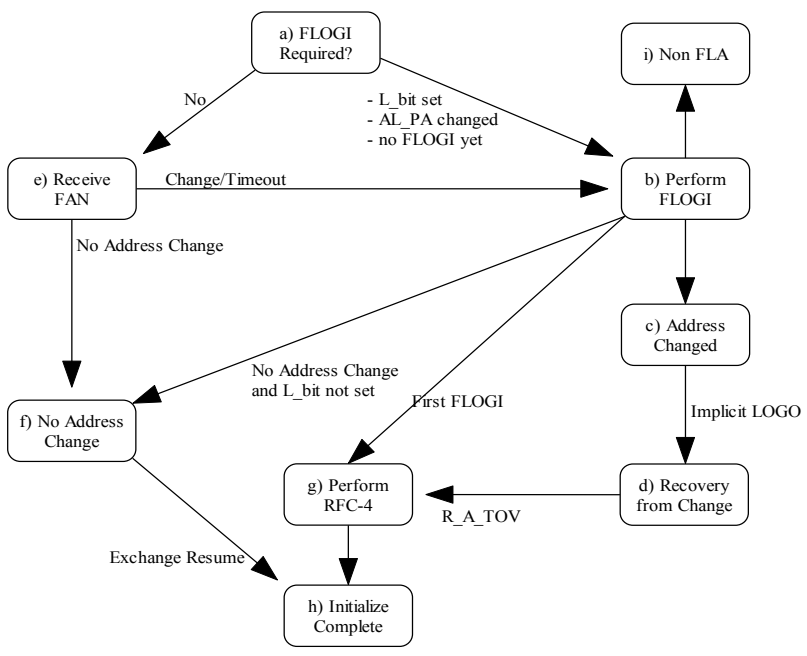
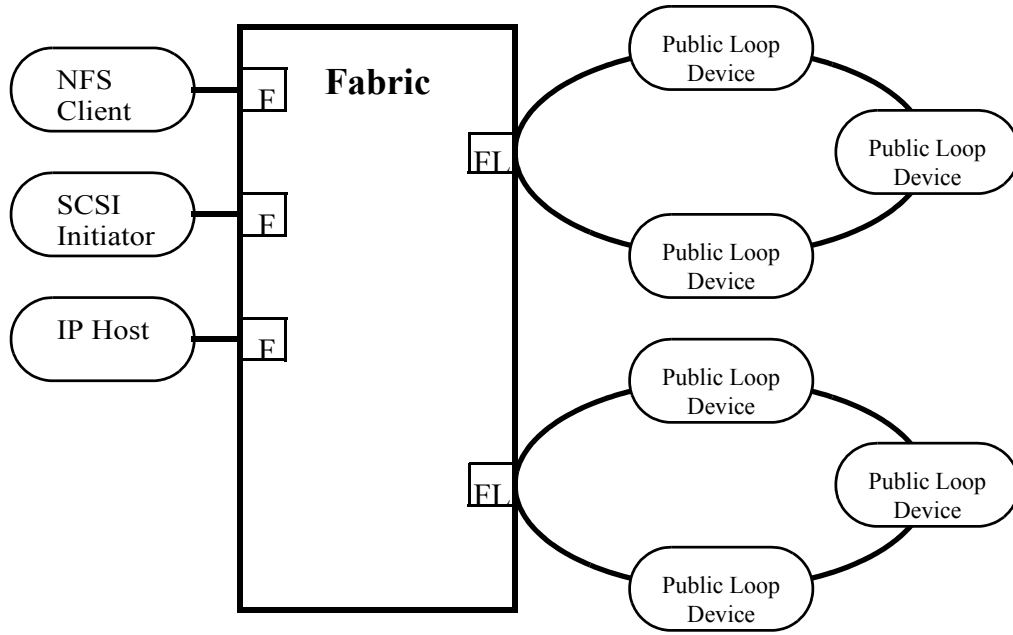
19.5 Public Loop Operation

The drive is designed to perform as a Public NL_Port described in the FLA 2.70 profile.

The drive will:

- tolerate unfairness of a FL_Port
- support 24bit addressing
- accept FAN extended link service
- perform FLOGI to Fabric F_Port (AL_PA = 00h, well known address 0FFFFFFEh)
- perform registration to a simple name server (AL_PA = 00h, well known address 0FFFFFFCh) using RFC-4 service.
- claim a fabric-assigned AL_PA during LIP in the LIFA sequence
- receive frames from a FL_Port for more than one S_ID in a single loop tenancy
- open L_Port at AL_PA = 00h and send frames to remote port
- use values of E_D_TOV, R_A_TOV as specified by the fabric in FLOGI ACC

19.5.1 NL_Port Initialization



During LIP, if the port had an AL_PA prior to the LIP, and the port had completed FLOGI, then the port attempts to acquire its Fabric-assigned AL_PA during LIFA sequence.

a) FLOGI Required

Upon completion of LIP, the port will implicitly log out with the Fabric and perform FLOGI if one or more of the following is true:

- the L_bit was set to one in at least the LISA sequence during the LIP.
- the port did not acquire the AL_PA it had prior to the LIP.
- the port did not have an AL_PA prior to the LIP.
- the port had not completed FLOGI prior to the LIP

b) Perform FLOGI

The port attempts to send FLOGI to the FL_Port by opening AL_PA = 00h using Full-Duplex. Meanwhile the port ignores any FAN ELS and responds to ADISC or PLOGI request sequences with an LS_RJT reply sequence, with a Reason code of “Unable to perform Command Request at this time”. If the FLOGI request fails for any reason, the port proceeds to step i). If no address change occurred and the L_bit is not set then the port proceeds to step f). If this is the first FLOGI or if the address has not changed and the L_bit is set, then the port proceeds to step g).

c) Address has changed

The port has determined that its own addressing information and/or that of the FL_Port has changed, or that it did not have completed addressing information (via FLOGI) prior to the initialization. The port discards all pending Exchanges, and performs implicit LOGO with all other ports. The port proceeds to step d).

d) Recovery from address change

The port waits R_A_TOV. During this time the port discards all frame received except for the PLOGI and ADISC request sequences and LOGO ACC reply sequences. The port responds to all PLOGI request with an LS_RJT reply sequence, with a Reason code of “Unable to perform Command Request at this time”. The port responds to each ADISC request sequence with a LOGO request sequence. The port proceeds to step g).

e) Receive FAN

If FLOGI is not required, the port waits E_D_TOV to receive FAN ELS. The port discards all frames received except ADISC and PLOGI requests until FAN is received. The port responds to ADISC and PLOGI request sequences with an LS_RJT reply sequence, with a Reason code of “Unable to perform Command Request at this time”.

After receiving the FAN ELS, if the port determines the FL_Port has the same address, F_Port_Name and Fabric_Name that the FL_Port had before initialization, the port proceeds to f). Otherwise, the port implicitly logs out with the Fabric and proceeds to b).

If the port does not receive a FAN ELS within E_D_TOV the port proceeds to b).

f) Address has not changed

The port resumes all suspended Exchanges with Remote Ports and with Local Ports that are known to be Public NL_Ports. If the port has any suspended Exchanges with Private Loop devices on the Local Loop, the port requires the private loop authentication described in PLDA profile. The port proceeds to step h).

g) Perform RFC-4

The port performs PLOGI with the Directory Server (AL_PA = 00h, well known address 0FFFFFFCh) and attempts an RFC-4 request with the Name Service. Regardless of this attempt of RFC-4, the port proceeds to step h).

h) Initialize complete

This completes initialization for the port.

i) Non FLA

The port is not connected to a loop that contains an FL_Port. The port completes initialization based on the rules described in PLDA profile.

19.6 SCSI Protocol

There are various operating conditions that prevent the Target from executing a SCSI command. This section describes each of these operating conditions and their relative priority.

19.6.1 Priority of SCSI Status Byte Reporting

After establishing the I_T_L nexus or I_T_L_Q nexus the Target must first determine whether command execution is allowed. Execution is deferred until a later time if the command must be added to the command queue. Execution may also be prevented by an internal Target condition that requires the reporting of a Check Condition, Queue Full, Busy, or Reservation Conflict Status. There are several different internal conditions to be active at the same time. The order in which the Target checks for each of these conditions determines their priority (highest priority first) as follows:

1. Check Condition status for invalid Logical Unit Number. (See Section 19.6.2, “Invalid LUN Processing” on page 279)
2. Check Condition status for Incorrect Initiator Connection.
3. Check Condition status for Unit Attention Condition (See Section 19.6.4, “Unit Attention Condition” on page 281)
4. Busy Status or Queue Full Status (See 19.6.3, “Command Processing During Execution of Active I/O Process” on page 279)
5. Check Condition status for Deferred Error Condition (See Section 19.6.7, “Deferred Error Condition” on page 283)
6. Check Condition status during Startup and Format operations (See Section “19.6.5, “Command Processing During Startup and Format Operations” on page 283)
7. Reservation Conflict status (See Section 19.6.9, “Command Processing while Reserved” on page 290)
8. Check Condition status for invalid command opcode
9. Check Condition status for invalid command descriptor block

The highest priority internal condition that prevents command execution is reported by the Target provided there is no bus error.

For all Check Conditions Sense data is built by the target provided a valid LUN address is known. Sense data is cleared by the Target upon receipt of any subsequent command to the LUN from the initiator receiving the Check Condition.

19.6.2 Invalid LUN Processing

Any value other than zero in the FCP_LUN field of the FCP_CMD IU is invalid.

The target's response to an invalid LUN varies with the command, as follows:

Inquiry: Execute the command, return the INQUIRY data that indicates unknown device type (byte 0 = 7Fh), and return GOOD status. All other bytes are valid (see 17.5, “INQUIRY (12)” on page 112).

Request Sense: Execute the command, return the sense data with the Sense Key set to Illegal Request and the Additional Sense Code and Additional Sense Code Qualifier set to LOGICAL UNIT NOT SUPPORTED, and return GOOD status (see also 17.34, “REQUEST SENSE (03)” on page 227).

All Others: Do not execute the command and return CHECK CONDITION status, along with the auto-sense data with the Sense Key set to Illegal Request and the Additional Sense Code and Additional Sense Code Qualifier set to LOGICAL UNIT NOT SUPPORTED.

In all cases, the target's response to the command for an invalid LUN does not affect the current execution of a command on the valid LUN for this initiator or any other initiator.

19.6.3 Command Processing During Execution of Active I/O Process

When the target is not executing any I/O processes, a new I/O process is permitted to execute (unless execution is prevented by another internal target condition listed in 19.6.1, “Priority of SCSI Status Byte Reporting” on page 279).

If an active I/O process exists when the target receives a new command, then the target determines if:

- the command is permitted to execute
- the command is added to the queue
- Queue Full status is to be returned
- Busy status is to be returned

If an active I/O process exists when the target receives a new command, then the target determines how the new command should be handled based on the following rules:

- Check Condition status is returned with sense key set to Logical Unit Not Ready if:
 - the startup operation or a format operation is active. See 19.6.5, “Command Processing During Startup and Format Operations” on page 283 for the exact conditions which cause this response.

Note: If a Unit Attention is pending when this condition exists, the sense key is set to Unit Attention rather than Logical Unit Not Ready since Unit Attention has a higher reporting priority (see 19.6.1, “Priority of SCSI Status Byte Reporting” on page 279).

- The command is permitted to execute if:
 - the command is a priority command (see 19.7, “Priority Commands” on page 290).
- the conditions to execute concurrently are met (see 19.10, “Concurrent I/O Process” on page 291).
- The command is added to the queue if:
 - any I/O process already exists at the target, and
 - this is not an incorrect initiator connection.
- Queue Full status is returned if:
 - the command would otherwise be added to the queue (according to the rules described above), but all slots in the queue are full, or
 - the command would otherwise be added to the queue (according to the rules described above), but all of the available queue slots not reserved for use by another initiator are full, or
 - a Format Unit command was previously queued but has not yet begun execution, or
 - the target is in a Degraded Mode (see 19.6.8, “Degraded Mode” on page 284) and a Start Unit command was previously queued but has not yet begun execution.
- Busy status is never returned.

19.6.4 Unit Attention Condition

The target generates a unit attention condition when one of the following occurs:

- The target has been reset

This includes a power-on reset or a reset caused by a Target Reset Task Management function or Reset LIP. In all of these cases, a unit attention condition is generated for each initiator. In addition, a process login (PRLI) will cause a Unit Attention Condition Power-On Reset for that initiator with an Additional Sense Code and Additional Sense Code Qualifier reported as Power-On Reset, Power-On Reset Occurred.

- MODE SELECT command has been executed

In this case, a unit attention condition is generated for all initiators except the one that issued the MODE SELECT command. The Additional Sense Code and Additional Sense Code Qualifier reported is MODE PARAMETERS CHANGED. The unit attention condition is generated if any of the current page parameters are set by the MODE SELECT command. The target does not check to see that the old parameters are different from the new parameters. For example: If the initiator issues a MODE SENSE command with a page code to report the current values followed by a MODE SELECT command with the same parameter list, a unit attention condition is generated despite the fact that the current parameters were not changed from their previous value. However, if the target detects an illegal parameter or error condition prior to modifying the current parameters, a unit attention condition is not generated since the parameters were not set. The unit attention condition is also not generated if the MODE SELECT command parameter list does not include any pages and only the header or header/block descriptor is present.

- FORMAT UNIT command has been executed

In this case, a unit attention condition is generated for all initiators except the one that issued the FORMAT UNIT command. The Additional Sense Code and Additional Sense Code Qualifier reported is NOT READY TO READY TRANSITION, (MEDIUM MAY HAVE CHANGED). This indicates that the block descriptor parameters from the last MODE SELECT command have been used and are now considered current values.

- WRITE BUFFER command to download microcode has been executed

In this case, a unit attention condition is generated for all initiators except the one that issued the WRITE BUFFER command. The Additional Sense Code and Additional Sense Code Qualifier reported is MICROCODE HAS BEEN CHANGED.

- Commands Cleared by another initiator

This unit attention condition is generated after an initiator sends a Clear Task Set Task Management function. The unit attention condition is generated for all other initiators with I/O processes that were either active or queued for the logical unit. The Additional Sense Code and Additional Sense Code Qualifier reported is COMMANDS CLEARED BY ANOTHER INITIATOR.

- LOG SELECT command with PCR bit has cleared parameters.

In this case, a unit attention condition is generated for all initiators except the one that issued the LOG SELECT command. The additional sense code and additional sense code qualifier reported is Log Select Parameters Changed.

- The registration or reservation made by a Persistent Reserve Out command was cleared by another initiator.

In this case, a unit attention condition is generated for the initiator that held the cleared registration or reservation.

- A Predictive Failure Analysis threshold has been reached and the Method of Reporting field of mode page 1Ch is 2h.

The unit attention condition persists for each initiator until that initiator clears the condition from the logical unit as described below. Several commands are handled as special cases during a unit attention condition. These cases are also discussed below.

If the target receives a command from an initiator before reporting a CHECK CONDITION status for a pending unit attention condition for that initiator, the target's response varies with the command as follows:

- Inquiry** Execute the command, return GOOD status, and preserve the unit attention condition.
- Report Luns** Same as above
- Request Sense** Execute the command, return any pending sense data, return GOOD status, and preserve the unit attention condition. If there is not any pending sense data, the sense data associated with the highest priority unit attention condition is returned and the highest priority unit attention condition is cleared for this initiator.
- All Others** Do not execute the command, return a CHECK CONDITION status, clear the highest priority unit attention condition for this initiator and return the associated sense data.

More than one unit attention condition may be generated for an initiator before that initiator clears the unit attention condition.

19.6.5 Command Processing During Startup and Format Operations

If the Target receives a command from an Initiator while the Target is executing a startup or format operation, the response of the Target varies with the command as follows:

INQUIRY	The drive sends inquiry data and returns appropriate status.
REQUEST SENSE	Executes the command, returns a Sense key of NOT READY and an Additional Sense Code of LOGICAL UNIT NOT READY and returns GOOD STATUS. The Additional Sense Code Qualifier that is returned depends on type of I/O processes that are active: For the START/UNIT STOP and the Auto-start operation, the qualifier returned is LOGICAL UNIT IS IN PROCESS OF BECOMING READY. For the FORMAT UNIT command, the qualifier returned is LOGICAL UNIT NOT READY, FORMAT IN PROGRESS, and the Sense key specific bytes are set to return the progress indication.
REPORT LUNS	The drive sends REPORT LUNS data and appropriate status.
ALL OTHER	The drive terminates the command with CHECK CONDITION status. The Sense data generated is described in Request Sense above.

19.6.6 Internal Error Condition

The Target generates an Internal Error condition for all Initiators when an internally initiated operation ends with an unrecoverable error.

An Internal Error condition causes Sense data to be generated and saved for all Initiators. The Error Code field of the Sense is set for a Current Error (70h) and the Sense Key is set to HARDWARE ERROR. Recovered errors are not reported.

The Internal Error condition persists for each Initiator until that Initiator clears the condition from the logical unit as described below. Several commands are handled as special cases during an Internal Error condition. These cases are also discussed.

If the Target receives a command from an Initiator while an Internal Error condition exists for that Initiator, the response of the Target varies with the command as follows:

INQUIRY	The drive executes the command with GOOD status and does not clear the Internal Error condition.
REQUEST SENSE	The drive executes the command, returns the sense data generated by the Internal Error condition, returns Good Status, and clears the Internal Error condition for that Initiator.
ALL OTHER	The drive terminates the command with a CHECK CONDITION status and clears the Internal Error condition.

19.6.7 Deferred Error Condition

Error code (71h) of sense data indicates that the Check Condition status returned is the result of an error or exception condition that occurred during execution of a previous command for which Good status has already been returned.

The drive creates an Deferred Error condition when

- Execution of a Format Unit command with the immediate bit of one ends with an error.

19.6.8 Degraded Mode

There are certain errors or conditions which may impair the ability of the drive to function normally. Rather than fail hard the drive is designed to be as responsive as possible. Also, in most cases, some action on the part of the initiator may be used to restore normal operation. This mode of limited operation is called Degraded Mode.

There are 3 conditions in the Degraded Mode:

- Media Degraded which is caused by one of the following conditions:
 - Context Load was started (by POR or Unit Start command) and the Target is under Self Configuration
 - Context Load Failure (Disable Logical)
 - Unit Stop command was issued after the Target successfully completed the Self Configuration
- Self Configuration Failure Degraded which is caused by one of the following conditions:
 - RAM Code, Read Failure
 - RAM Code, Revision Mismatch
- Format Command Failure Degraded. This condition is caused when Format Unit command failed or was interrupted abnormally (Mode Page 0, byte 5, bit 4 FDD controls Format Degraded mode)

19.6.8.1 Response to SCSI Command in Degraded Mode - Becoming Ready

Table 267: Media Degraded Mode Becoming ready

Command (w/Option)	Response
Request Sense	Executed. The Target may return Sense Key 02h (Not Ready) ASC/ASCQ 0401h (In Process of Becoming Ready)
Inquiry (EVPD=0)	Executed
Inquiry (EVPD=1)	Executed and Check Condition is returned with Sense Key 05h (Illegal Request) ASC/ASCQ 2400h (Invalid Field in CDB)
Test Unit Ready	Executed and Check Condition is returned with Sense Key 02h (Not Ready) ASC/ASCQ 0401h (In Process of Becoming Ready)
Start Stop Unit (Start)	<p>Executed</p> <ul style="list-style-type: none"> - Success: Good Status is returned. Media Degraded Mode is cleared - Context Load Failure: Check Condition with Sense Key 04h (Hardware Error) ASC/ASCQ 4400h (Internal Target Error) - Self Configuration Failure: Check Condition with Sense Key 04h (Hardware Error) ASC/ASCQ 4080h (Diag Fail- Bring up Fail)
Other Commands	Not Executed. Check Condition Status is returned with Sense Key 02h (Not Ready) ASC/ASCQ 0401h (In Process of Becoming Ready)

19.6.8.2 Response to SCSI Command in Degraded Mode - Context Load

Table 268: Media Degraded Mode - Context Load Failure

Command (w/Option)	Response
Request Sense	Executed. The Target may return Sense Key 04h (Hardware Error) ASC/ASCQ 4400h (Internal Target Failure)
Inquiry (EVPD=0)	Executed
Inquiry (EVPD=1)	Executed and Check Condition is returned with Sense Key 05h (Illegal Request) ASC/ASCQ 2400h (Invalid Field in CDB)
Test Unit Ready	Executed and Check Condition is returned with Sense Key 04h (Hardware Error) ASC/ASCQ 4400h (Internal Target Failure)
Start Stop Unit (Start)	Executed - Success: Good Status is returned. Media Degraded Mode is cleared - Context Load Failure: Check Condition with Sense Key 04h (Hardware Error) ASC/ASCQ 4400h (Internal Target Failure) - Self Configuration Failure: Check Condition with Sense Key 04h (Hardware Error) ASC/ASCQ 4080h (Diag Fail- Bring up Fail)
Start Stop Unit (Stop)	Executed. Good Status is returned. Media Degraded Mode is NOT cleared
Other Commands	Not Executed. Check Condition Status is returned with Sense Key 04h (Hardware Error) ASC/ASCQ 4400h (Internal Target Failure)

19.6.8.3 Response to SCSI Command in Degraded Mode - Drive issued received unit stop Command

Table 269: Media Degraded Mode - Spindle Stopped by Unit Stop Command

Command (w/Option)	Response
Request Sense	Executed. The Target may return Sense Key 02h (Not Ready) ASC/ASCQ 0402h (Initialize Command Required)
Inquiry (EVPD=0)	Executed
Inquiry (EVPD=1)	Executed
Test Unit Ready	Executed and Check Condition is returned with Sense Key 02h (Not Ready) ASC/ASCQ 0402h (Initialize Command Required)
Start Stop Unit (Start)	<p>Executed</p> <ul style="list-style-type: none"> - Success: Good Status is returned. Media Degraded Mode is cleared - Context Load Failure: Check Failure: Check Condition with Sense Key 04h (Hardware Error) ASC/ASCQ 4400h (Internal Target Failure) - Self Configuration Failure: Check Condition with Sense Key 04h (Hardware Error) ASC/ASCQ 4080h (Diag Fail- Bring up Fail)
Start Stop Unit (Stop)	Executed. Good Status is returned. Media Degraded Mode is NOT cleared
Other Commands	Not Executed. Check Condition Status is returned with Sense Key 04h (Hardware Error:) ASC/ASCQ 4400h (Internal Target Failure)

19.6.8.4 Self Configuration Failure Degraded Mode

Table 270: Self Configuration Failure Degraded Mode

Command (w/Option)	Response
Request Sense	Executed. The Target may return Sense Key 04h (Hardware Error) ASC/ASCQ 4080h (Diag Fail- Bring up Fail)
Inquiry (EVPD=0)	Executed
Inquiry (EVPD=1)	Executed and Check Condition is returned with Sense Key 05h (Illegal Request) ASC/ASCQ 2400h (Invalid Field in CDB)
Test Unit Ready	Executed and Check Condition is returned with Sense Key 04h (Hardware Error) ASC/ASCQ 4080h (Diag Fail- Bring up Fail)
Start Stop Unit (Start)	Executed - Success: Good Status is returned. Media Degraded Mode is cleared - Context Load Failure: Check Condition with Sense Key 04h (Hardware Error) ASC/ASCQ 4400h (Internal Target Failure) - Self Configuration Failure: Check Condition with Sense Key 04h (Hardware Error) ASC/ASCQ 4080h (Diag Fail- Bring up Fail)
Write Buffer (Download and Save)	Executed. - Success: Good Status is returned. Media Degraded Mode is cleared - Self Configuration Failure: Check Condition with Sense Key 04h (Hardware Error) ASC/ASCQ 4080h (Diag Fail- Bring up Fail) Sense Key 02h (Not Ready) ASC/ASCQ 4085h (Diag Fail-RAM code NOT load)
Other Commands	Not Executed. Check Condition Status is returned with Sense Key 04h (Hardware Error) ASC/ASCQ 4080h (Diag Fail- Bring up Fail)

19.6.8.5 Format Command Failure Degraded Mode

Table 271: Format Command Failure Degraded Mode

Command (w/Option)	Response
Request Sense	Executed. The Target may return Sense Key 02h (Not Ready) ASC/ASCQ 3100h (Format Corrupted) Sense Key 03h (Medium Error) ASC/ASCQ 3100h (Format Corrupted)
Inquiry (EVPD=0)	Executed
Inquiry (EVPD=1)	Executed
Test Unit Ready	Executed and Check Condition is returned with Sense Key 02h (Not Ready) ASC/ASCQ 3100h (Format Corrupted)
Format Unit	Executed - Success: Good Status is returned. Format Degraded Mode is cleared - Failure: Check Condition Status is returned and Format Degraded Mode is NOT cleared.
Other Commands	Not Executed. Check Condition Status is returned with Sense Key 03h (Medium Error) ASC/ASCQ 3100h (Format Corrupted)

Note: Mode Page 0 byte 5 bit 4 (FDD) = 0

19.6.9 Command Processing while Reserved

A logical unit is reserved after successful execution of the Reserve command. Each time a Reserve command is executed successfully, the Target records the SCSI ID of the Initiator that made the reservation and the SCSI ID of the Initiator that is to receive the reservation. This information is needed to determine whether subsequent commands should be permitted or if the Reservation Conflict Status should be reported. The Initiator that made the reservation is the Initiator that issued the Reserve command. The Initiator to receive the reservation may be either the same or a different Initiator (third-party reservation).

If the logical unit is reserved when a new command is received, the Target examines the command opcode and the SCSI ID of the issuing Initiator to determine whether a Reservation Conflict Status should be returned based on the following rules:

If the issuing Initiator is the one that made the reservation and also the one to receive the reservation, then all commands are permitted.

If the issuing Initiator is neither the one that made the reservation nor the one to receive the reservation, then

- A Request Sense or Inquiry command is permitted.
- A Release command is permitted but is ignored.
- Any other command results in a Reservation Conflict Status.

If the issuing Initiator is the one that made the reservation but is not the one to receive the reservation, then

- An Inquiry, Request Sense, Reserve, or Release command is permitted.
- Any other command results in a Reservation Conflict Status.

If the issuing Initiator is not the one that made the reservation but is the one to receive the reservation, then

- A Reserve command results in a Reservation Conflict Status.
- A Release command is permitted but is ignored.
- Any other command is permitted.

If a Reservation Conflict Status is not reported and the command is permitted, then the Target checks the next highest priority internal condition to determine whether execution is allowed. See Section 19.6.1, “Priority of SCSI Status Byte Reporting” on page 279.

19.7 Priority Commands

Certain SCSI commands always execute without returning a Busy Status or Reservation Conflict Status in response to the command. These commands are

- Inquiry
- Request Sense
- Report LUNs
- Test Unit Ready

These commands are executed prior to attempting to complete the execution of any other pending command in the queue. These commands are never queued.

19.8 Command Queuing

When the initiator specifies that the drive shall disable command queuing, the initiator must send only untagged commands. When the initiator specifies that the target shall enable command queuing, the initiator may send either tagged or untagged command, but shall not use both at the same time.

The following commands are never queued.

- Priority Commands (i.e.: Request Sense and Inquiry)
- Commands for an invalid LUN.

19.8.1 Queue Depth

Any initiator can queue at least one command at any time irrespective of the actions of any other initiators in the system. A single initiator may queue up to 128 commands, if no other initiator has more than one command in the queue, although at times this maximum may be reduced as the drive can reserve command blocks for internal use.

19.8.2 Queue Full Status

The drive will respond with QUEUE FULL status to a SCSI command when all queue slots are utilized. The SCSI command is not placed in the command queue under this condition.

19.8.3 Termination of I/O Processes

Normal termination of I/O processes occurs when the target returns SCSI status. I/O processes may also be terminated by the following:

- An ABORT TASK terminates the specified I/O process from the issuing initiator
- An ABORT TASK SET terminates all I/O processes from the issuing initiator
- A CLEAR TASK SET, TARGET RESET or reset terminates all I/O processes from all initiators
- Failure by an initiator to initiate Exchange Authentication within RR_TOV of the completion of Loop Initialization results in an implicit logout and termination of all associated I/O processes (21.8.5 Effects of LIP on Command Queuing on a page 291)

19.9 Command Reordering

Command reordering is supported when enabled by the Queue Algorithm Modifier in mode page 0A (see 17.10.9, “Mode Page 0A (Control Mode Page Parameters)” on page 169.

19.10 Concurrent I/O Process

Concurrent command are always allowed to execute concurrently with non-priority commands. A second priority command received while a priority command is being executed is put at the head of the command queue.

- WRITE commands when another WRITE command is an active I/O process
- READ commands when another READ command is an active I/O process

When a concurrent command ends in CHECK CONDITION status, the QErr bit on the Mode Page 0Ah will determine how other active I/O processes from the same initiator for that drive will be handled.

19.11 Write Cache

If the WCE (Write cache enable) bit is 1, the drive returns Good Status and Task complete message and goes to Bus Free immediately after receiving the data of the last sector before actually writing the data onto the media.

If the drive detects an error after it returns a Good Status, the drive sets a Deferred Error (Error Code of sense data = 71h) and a following command will be returned with Check Condition and the Contingent allegiance condition is established. Under the Contingent allegiance condition all queued processes including commands from other initiators are suspended.

19.12 Automatic Rewrite/Reallocate

The target supports auto reallocation for all media access commands. Auto reallocation cannot be disabled, and the ARRE setting is ignored.

19.13 Segmented Caching

19.13.1 Overview

Segmented Caching divides the data buffer into several smaller buffers. Each buffer is used as Read/ Write/Read-Ahead buffer.

19.13.2 Read Ahead

The Read Ahead function consists of reading data that the Initiator has not yet requested to the drive buffer. This function is intended to improve performance for an initiator that frequently accesses sequential data with successive SCSI read commands.

The drive initiates the Read ahead function when the following conditions exist:

- Read, Verify and Write and Verify is received.
- The consecutive LBA of the requested LBA is not available in the buffer.

Even if an error occurs during the Read ahead, the error will not be reported to the Initiator. The data read before the error occurred will be stored as valid data by the Read Ahead function.

19.14 Multiple Initiator Systems

This section describes how the target behaves in a multiple initiator system. Up to 64 initiators may be supported at any one time.

19.14.1 Sense Data

A separate sense data area is reserved for each initiator. Each area is maintained independently. This allows a command from one initiator to complete with a CHECK CONDITION status and generate sense data without being affected by a subsequent command from a different initiator. There is no requirement for the first initiator to send a REQUEST SENSE command to retrieve the Sense Data prior to the execution of a command from a different initiator.

19.14.2 Mode Pages

A single set of Mode pages is maintained. This includes both current and saved parameters. If a MODE SELECT command is executed that updates the current parameters, a unit attention condition is generated for all initiators except the one that issued the command. See 19.6.4, "Unit Attention Condition" on page 281 for more information.

19.15 Enclosure Services

Enclosure Services allow the Host Computer System to access information held by the Enclosure Micro controller using the same industry standard interface as the enclosed devices.

This drive supports the ESI (Enclosure Services Interface) for this purpose, as described in the Functional Specification.

The drive sets the EncServ bit to 1b in the Standard INQUIRY Data indicating that it supports enclosure services. This indicates that the Host may use the SEND DIAGNOSTIC and RECEIVE DIAGNOSTIC RESULTS commands to access information held by the Enclosure. This information is in the form of 'Diagnostic Pages'.

Diagnostic Pages 01h to 0Fh are reserved for Enclosure Services.

There are two distinct forms of Enclosure Services. The EncServ bit does not specify which form the Target and Enclosure are capable of (if any). The Host can discover this by using a RECEIVE DIAGNOSTIC RESULTS command to request Diagnostic Page 00h.

When requested to communicate with the enclosure via the SEND and RECEIVE DIAGNOSTIC commands, the drive goes through a 'discovery phase' with the enclosure as defined in the SFF-8067 standard. The results of this discovery phase determine whether the enclosure: 1) does not support Enclosure Services, 2) supports only the 'short' Enclosure Services mode as defined by SFF-8045, or 3) supports the 'long' Enclosure Services mode as defined by SFF-8067. After the discovery phase,

the drive behaves in accordance with the Enclosure Services mode capabilities of the enclosure. The drive does not attempt to remember the Enclosure Services capabilities of the enclosure, but rather, in accordance with the SFF-8067 standard, re-performs the discovery phase for each SEND and RECIEVE DIAGNOSTIC command received.

Short

This mode is defined for use when an Enclosure simply supplies an 7-bit parameter as defined by the SFF-8045 standard to the drive.

When mounted in this type of Enclosure, an attempt to access any Enclosure Service Diagnostic Page using the SEND DIAGNOSTIC command will terminate the command with CHECK CONDITION status. The sense data will have a Sense Key of Illegal Request and an Additional Sense Code of UNSUPPORTED ENCLOSURE FUNCTION.

Requesting any Enclosure Service Diagnostic Page using the RECEIVE DIAGNOSTIC RESULTS command will return only Diagnostic Page 08h (Short Enclosure Status).

Long

This form is defined for use when the Enclosure and target are able to communicate using the secondary interface as defined in the SFF-8067 standard.

When mounted in this type of Enclosure, any attempt to access an Enclosure Service Diagnostic Page using the SEND DIAGNOSTIC or RECEIVE DIAGNOSTIC RESULTS commands results in the Diagnostic Page being transferred to/from the enclosure. CHECK CONDITION status and sense data with an appropriate Sense Key and Additional Sense Code is generated if the transfer fails or cannot be performed.

Note: For a more detailed description of Enclosure Services, please refer to the following specifications:

- SFF-8067 Specification

19.15.1 Enclosure Initiated ESI

This drive supports Enclosure Initiated ESI as described in the SFF-8067 standard.

19.16 Multiple Initiator Environment

19.16.1 Initiator Sense Data

Separate sense data is reserved for each I-T-L. Each sense data is maintained independent of commands from other initiators.

19.16.2 Initiator Mode Select/Mode Sense Parameters

A single shared copy of the Mode Select/Mode Sense parameters is maintained by the drive. This includes both the current and saved parameters.

19.17 Reset

Reset actions will return the drive to a known, initialized state.

This device supports the Hard reset option as defined in the SCSI standards (see 7.1.9 Dev_Ctrl_Code_x on page 33) and the reset sources discussed below.

19.17.1 Reset Sources

There are four sources of resets detected by the target:

Reset Name	Reset Source
Power-On Reset	This is the signal generated by the hardware at initial power-on
Self-Initiated reset	This is a software-generated reset that occurs when a catastrophic error is detected by the microcode.
LIP Reset	This is a LIP(AL_PS_PD) primitive sequence where the AL_PD matches the AL_PA of the drive.
Target Reset	This is an FCP_CMD IU with the TARGET RESET TMF flag set.
Reset LUN	This is a FCP_CMD IU with the RESET LUN TMF flag set.

19.17.2 Reset Actions

The action taken by the drive following a reset is dependent on the source of the reset.

19.17.2.1 Power-On reset and Self-Initiated reset

These two reset conditions cause the following to be performed in the order shown:

- A power-up sequence
- A startup sequence is necessary to put the drive in a ready state

These reset conditions cause the following actions:

- If the reset occurs during the power-up sequence, the power-up sequence is re-started.
- If a start-up sequence has not yet completed, the start-up sequence is restarted. Note: The power-up sequence is not re-run, since it has already completed.
- If the reset occurs while a physical sector is being written, the WRITE operation is disabled at the end of the current sector. The media is not corrupted.

19.18 Diagnostics

The drive will execute Power on Diagnostics at power on time to assure the correct operation of the drive by validating components (ROM, RAM, Sector Buffer, EEPROM, HDC, PLI Capacitor), checking stored information in the Reserved Area and EEPROM, and verifying fault detects circuits.

Self-test can be invoked by issuing a SEND DIAGNOSTIC command.

19.18.1 Power on Diagnostics

At power on time the following tests are executed:

1. Validation of ROM and EEPROM
2. RAM test for internal RAM
3. Test and Initialize HDC registers
4. RAM test for Sector Buffer
5. PLI selftest
6. Validation of RAM code and data table (RDM, Log, Mode Page) from the Reserved Area

Faults detected before successful completion of the HDC section could prevent the drive from responding to a selection.

Faults detected after the successful completion of the HDC test section will be reported as CHECK CONDITION status to the Initiator on the first command issued after a fault is detected (except for the INQUIRY, REPORT LUNS and REQUEST SENSE commands). The INQUIRY, REPORT LUNS and REQUEST SENSE commands will always be responded with a GOOD status. Detecting a fault during power on will not terminate execution of the tests nor will it terminate the power on process.

19.18.2 Self-test via SEND DIAGNOSTIC Command

19.18.2.1 Default Self-test

The default self-test is invoked by the SlfTst bit in the SEND DIAGNOSTIC command. The response is simply a GOOD status if the test is successful or a CHECK CONDITION status if the test fails.

The following tests are performed by the default self-test (in the order defined):

1. **Logical Enable Check to determine if the media is accessible logically.**
2. **Write, Read and Compare test** is a drive read/write test. It writes data to a predefined location in the reserved area and then reads it back and validates the content.

19.18.2.2 Short and Extended Self-tests

There are two other types of self-tests that may be invoked using the Function Code field in the SEND DIAGNOSTIC command: a short self-test and an extended self-test. The tests performed in the short and extended self-tests are described later. The time required by a logical unit to complete its extended self-test is specified in the Extended self-test Completion Time field in the Control Mode Page. The results of self-test can be retrieved via the LOG SENSE command for Log Page 10.

19.18.2.3 Self-test Modes

There are two modes for short and extended self-tests: a foreground mode and a background mode. These modes are described in the following clauses.

Foreground mode

When the drive receives a SEND DIAGNOSTIC command specifying a self-test to be performed in the foreground mode, the drive will return status for that command after the self-test has been completed. While performing a self-test in the foreground mode, the drive will respond to all commands except INQUIRY, REPORT LUNS, and REQUEST SENSE with a CHECK CONDITION status, a sense key of NOT READY and an additional sense code of LOGICAL UNIT NOT READY - SELF-TEST IN PROGRESS.

If the drive is performing a self-test in the foreground mode and a test error occurs, the drive will update the self-test results log page and report CHECK CONDITION status with a sense key of HARDWARE ERROR and an additional sense code of LOGICAL UNIT FAILED SELF-TEST. The application client may obtain additional information about the failure by reading the self-test results log page.

An application client may terminate a self-test that is being performed in the foreground mode using an ABORT TASK, ABORT TASK SET, or CLEAR TASK SET task management function. If the drive receives an ABORT TASK, ABORT TASK SET, or CLEAR TASK SET task management function while performing a self-test in the foreground mode, it will abort the self-test and update the self-test results log page.

Background mode

When the drive receives a SEND DIAGNOSTIC command specifying a self-test to be performed in the background mode, the drive will return status for that command as soon as the command descriptor block has been validated. After returning status for the SEND DIAGNOSTIC command specifying a self-test to be performed in the background mode, the drive will initialize the self-test results log page as follows. The Function Code from the SEND DIAGNOSTIC command will be placed in the Function Code field in the log page. The self-test Results field shall be set to 0Fh. After the self-test results log page is initialized, the drive will begin the first self-test segment.

While the device server is performing a self-test in the background mode, it shall terminate with a CHECK CONDITION status any SEND DIAGNOSTIC command it receives that meets one of the following criteria:

- a. The SlfTst bit is one
- b. The Function Code field contains a value other than 000b or 100b.

When terminating the SEND DIAGNOSTIC command, the sense key shall be set to NOT READY and the additional sense code shall be set to LOGICAL UNIT NOT READY, SELF-TEST IN PROGRESS. While performing a self-test in the background mode, the drive will suspend the self-test to service any other command other than SEND DIAGNOSTIC (with Function Code field set to 100b) WRITE BUFFER (with the mode set to any download microcode option), FORMAT UNIT and START UNIT STOP command. Suspension of the self-test to service the command will occur within 2 seconds. If SEND DIAGNOSTIC (with Function Code field set to 100b), WRITE BUFFER (with the mode set to any download microcode option), FORMAT UNIT or START UNIT STOP command is received, the drive will abort the self-test, update the self-test log, and service the command within two seconds after the command descriptor block has been validated.

An application client may terminate a self-test that is being performed in the background mode by issuing a SEND DIAGNOSTIC command with the Function Code field set to 100b (Abort background self-test function).

Elements common to foreground and background self-test modes

The Progress Indication field returned in response to a REQUEST SENSE command may be used by the application client at any time during execution of a self-test to poll the progress of the test. While executing a self-test unless an error has occurred, the drive will respond to a REQUEST SENSE command by returning a sense key of NOT READY and an additional sense code of LOGICAL UNIT NOT READY - SELF-TEST IN PROGRESS with the sense key specific bytes set for progress indication.

The application client may obtain information about the twenty most recently completed self-tests by reading the self-test results log page. This is the only method for an application client to obtain information about self-tests performed in the background mode. The default self-test results are not logged in the log page.

Tests performed in the Short and Extended Self-test

The following table defines the tests performed in the short and extended self test. They are defined by their segment number which is also used to report Self-Test Results, in Log Sense Page 10. Note that the only difference between the Short and the Extended tests, is the sequential verify test in segment 7h. Also note that either of these tests can be run in foreground or back-

ground mode as previously described.

Table 272: Short and Extended Self-Test Description

Segment Number	Short Self-Test	Extended Self-Test	Test Description
1h	Drive Ready Test		Internal check to insure drive is “ready”, similar to a Test Unit Ready command.
2h	Drive Diagnostics		This test is comprised of the Default Self Test as defined in Section 19.18.2.1, “Default Self-test” on page 296
3h	PLI Test		Drain and recharge Caps
4h	SMART		Perform SMART testing and check results to ensure that SMART threshold criteria are not exceeded
5h	Low Level Format check		Check to insure that the media is currently not in the MEDIA FORMAT CORRUPTED state.
6h	Reserved Area check		Write/Read test in a predefined location in the drive's Reserved Area of the drive.
7h	- Verify First 300MB - Verify Last 100 MB	Verify all LBA's	Sequential verify operation. Ensure that no uncorrectable errors occur within the verify range.
8h	Recheck SMART		Same as segment 4h.

19.18.2.4 Background Medium Scan

For a related function, see Mode Page 1C (Informational Exceptions Control), page 175

19.19 Idle Time Function

The drive periodically saves data in logs and S.M.A.R.T. counters in the reserved area of the drives. The information is used by the drive to support various SCSI commands and for the purpose of failure analysis.

19.20 Command Time out Limits

The 'Command Time-out Limits' are defined as the time period from the SCSI Arbitration phase through the SCSI Task complete message, associated with a particular command.

The following times are for environments where Automatic Reallocation is disabled and there are no queued commands.

19.20.1 Reassignment Time

The drive should be allowed a minimum of 5 seconds to complete a "Reassign Blocks" command.

19.20.2 Format Time

Approximately 3 minutes should be allowed to complete a "Format Unit" command.

19.20.3 Start/Stop Unit Time

The drive should be allowed a minimum of 10 seconds to complete a "Start Stop Unit" command (with Immed bit = 0). Initiators should also use this time to allow startup sequences initiated by auto start ups and "Start Stop Unit" commands (with Immed bit = 1) to complete and place the drive in a "ready for use" state.

19.20.4 Medium Access Command Time

The time-out limit for medium access commands that transfer user data or non-user data or both should be a minimum of 30 seconds. These commands are

- Pre-Fetch
- Read
- Read Defect Data
- Seek
- Send Diagnostic (Function Code = 0)
- Read Long
- Reassign Blocks
- Write
- Write and Verify
- Write Buffer
- Write Long
- Write Same
- Verify

Note: The 5-second limit assumes the absence of bus contention and data transfers of 64 blocks or less. This time should be adjusted for anticipated bus contention and if longer user data transfers are requested.

19.20.5 Time-out Limits for Other Commands

The drive should be allowed a minimum of 5 seconds to complete these commands:

- Inquiry
- Log Select
- Log Sense
- Mode Select
- Mode Sense
- Persistent Reserve In/Out
- Read Buffer
- Read Capacity
- Release
- Request Sense
- Reserve
- Set/Report Device Identifier
- Start/Stop Unit (with Immed bit = 1)
- Synchronize Cache
- Test Unit Ready

The command time-out for a command that is not located at the head of the command queue should be increased by the sum of command time-outs for all of the commands that are performed before it is.

19.21 Recommended Initiator ERP

The Drive's design points for error reporting to the system assumes certain system action for the error return codes. These assumptions are:

1. SCSI protocol will be the first priority in reporting errors.
2. The system will maintain a log of all reported errors.
3. System architecture should include all error handling recommendations made in this section. Deviations should have mutual agreement between Drive development and system integration.

This section is directed toward documenting the assumptions made by the Drive that the system is expected to implement. The two error classes that the system should be concerned with are DATA and NON-DATA errors.

Data errors are those errors that deal with the handling of data to and from the MEDIA and are identified by the Additional Sense Code contained in the sense data. The Additional Sense Codes for data errors are as follows:

- 11 - Unrecovered read error
- 17 - Recovered read error

Nondata errors are those errors that do not have a direct relationship with transferring data to and from the media. Nondata errors can include data handling if the media is not associated with the error (that is, interface error).

The system action assumed for each class of error is outlined here.

19.21.1 Drive Service Strategy

The Drive service strategy is defined so the customer will be able to use the system as soon after a failure is detected as possible. The first priority is to replace the entire drive to make the system operational with minimal service time. The service representative should:

1. Back up all the customer data on this drive if possible
2. Replace the complete drive
3. Restore the customer data
4. Return the drive to customer service

19.21.2 Recommendations for System Error Log

The system error log should contain information about the Drive error that will allow recovery actions. The system error logs should contain all the error information returned in the sense data. At a minimum, the following information about each error occurrence should be logged.

- Valid bit and error code (Sense byte 0)
- Sense Key (Sense byte 2)
- Information bytes (Sense bytes 3 through 6)
- Command specific information (Sense bytes 8 through 11)
- Additional Sense Code (Sense byte 12)
- Additional Sense Code Qualifier (Sense byte 13)
- Field Replaceable Unit (Sense byte 14)
- Sense Key Specific (Sense bytes 15, 16, and 17)
- Vendor Unique error information (Sense bytes 20 through 23)

19.21.3 Data Recovery Procedure

No action can be taken on hard or soft read errors. Block retirement happens automatically based on the block retirement policy in the firmware. LBAs that report a hard read error will become readable after a write. Until a write command is received for the affected LBAs, a hard error will be reported on a read to the affected LBAs.

19.21.4 Nondata Error Recovery Procedure

The Drive will follow a logical recovery procedure for nondata errors. The initiator options for non-data errors are limited to logging the error, retrying the failing command, or replacing the drive.

These recovery procedures assume the initiator practices data back-up and logs errors at the system level for interrogation by service personnel.

19.21.4.1 Drive Busy

The Drive is busy performing an operation. **This is not an error condition.** The initiator can test for completion of the operation by issuing *Test Unit Ready (00)* (or media access) commands.

- If the *Test Unit Ready (00)* (or media access) command completes with *Check Condition Status* then issue a *Request Sense (03)*
 - If the specified recovery procedure for the sense data is for a condition other than drive busy, follow the recovery procedure for the condition reported.
 - If the specified recovery procedure for the sense data is for a drive busy condition, then continue re-issuing the *Test Unit Ready (00)* and *Request Sense* commands for the duration of a media access time-out or until the drive returns *Good Status*.
 - If the drive has been busy for longer than the limit specified in Section 19.20, “Command Time out Limits” on page 299, then service the drive using the service guidelines recommended in Section 19.21.1, “Drive Service Strategy” on page 301. Otherwise return to normal processing.
- If the *Test Unit Ready (00)* (or media access) command completes with *Good Status*, then return to normal processing.

19.21.4.2 Unrecovered Drive Error

The initiator should retry the failing command.

5. If the retry of the failing command completes with *Good Status* or recovered Sense Key, follow the recovery procedure in Section 19.21.4.3, “Recovered Drive Error” on page 303.
6. If the retry of the failing command completes with hardware error sense, verify there is no outside cause (e.g., power supply) for the failure, then retry the failing command.
 - a. If the retry of the failing command completes with *Good Status*, follow the recovery procedure in next Section 19.21.4.3, “Recovered Drive Error” on page 303.
 - b. If the retry of the failing command completes with Recovered sense or Hardware error sense, then service the drive using the service guideline recommended in Section 19.21.1, “Drive Service Strategy” on page 301.

19.21.4.3 Recovered Drive Error

The Initiator should log the error as soft with the recovery level.

19.21.4.4 Drive Not Ready

The initiator should do the following:

1. Issue a *Start Stop Unit (1B)* command.
2. Verify that the drive comes ready within the time specified in Section Table 8: , “SSD Response time” on page 9.
3. If the drive fails to come ready within the specified time, service the drive using the service guidelines specified in Section 19.21.1, “Drive Service Strategy” on page 301.
4. Retry the failing command.
 - a. If the failing command completes with *Good Status*, log the error as recovered.
 - b. If the failing command completes with Not Ready sense, verify there is no outside cause (for example, the power supply). Then service the drive using the service guidelines specified in Section 19.21.1, “Drive Service Strategy” on page 301.

19.21.4.5 Degraded Mode

Refer to Section 19.6.8, “Degraded Mode” on page 284, for the definition of this state. There are three causes for entering degraded mode. In all cases the Sense Key is *Not Ready*. The causes are the following:

1. Sense Code/Qualifier of *Logical Unit Not Ready, initializing command required*. The media is not accessible. This may not be an error condition. The initiator should issue a *Unit start (1B)* command to enable media access. If the Drive fails to come ready in the time specified in Section 19.20, “Command Time out Limits” on page 299, service the drive using the service guideline recommended in Section 19.21.1, “Drive Service Strategy” on page 301.
2. Sense Code/Qualifier of *Diagnostic Failure*. Failure of a Send Diagnostic self test, a start up sequence, or other internal target failures.
 - Failure of a send diagnostic self test or a start up sequence.

This failure is the result of the diagnostics that are executed during power on or when the *Send Diagnostic (1D)* command is executed detecting a failure. As with the RAM code not loaded and the configuration data not loaded, the recovery is either a power cycle or issuing the *Send Diagnostic (1D)* command with the self test bit set active.

Recovery for a failed Send Diagnostic (1D) is achieved in one of the following ways:

Executing the Send Diagnostic (1D) command

Power cycling the drive

If the failure repeats, service the drive using the service guideline recommended in Section 19.21.1, “Drive Service Strategy” on page 301.

Recovery for a failed power up sequence is achieved in one of the following ways:

Issuing a Unit start (1B) command

Power cycling the drive.

If the failure repeats, service the drive using the service guideline recommended in Section 19.21.1, “Drive Service Strategy” on page 301.

- Internal target failures

Recovery of this condition is either a power cycle or successful completion of the Send Diagnostic (1D). Service the drive using the recommended service guidelines specified in Section 19.21.1, “Drive Service Strategy” on page 301, if the power cycle or the Send Diagnostic (1D) command fail to complete successfully.

3. Sense Code/Qualifier of **Format Command Failed** Format Unit (04).

Recovery from a failed Format Unit (04) is achieved by retrying the command. If the command fails a second time, service the drive following the procedure defined in Section 19.21.1, “Drive Service Strategy” on page 301.

If the above defined recovery procedures fail to clear the degraded mode condition, the Drive should be replaced. Follow the procedure in Section 19.21.1, “Drive Service Strategy” on page 301, when replacing the drive.

19.21.4.6 Interface Protocol

For all interface protocol errors, the initiator should complete the following steps:

1. Correct the parameter that caused the Illegal Request
2. Retry the failing command
3. If the first retry of the failing command completes with
 - *Good Status*, log the error as recovered
 - *Check Condition Status* with sense data for an Illegal Request, verify there is no outside cause (for example, the power supply) for the failure
 - *Other*, follow the recommendations for the error condition reported. Retry the failing command. If this retry of the failing command completes with
 - *Good Status*, log the error as recovered
 - *Check Condition Status* with sense data for an Illegal Request, service the drive using the service guideline recommended in Section 19.21.1, “Drive Service Strategy” on page 301.
 - *Other*, follow the recommendations for the error condition reported.

19.21.4.7 Aborted Command

The initiator should determine the cause from the Additional Sense Code (byte 12):

- Sense Key = B (Aborted Command) with Additional Sense Codes of 1B, 25, 43, 49, and 4E are initiator caused abort conditions. The initiator should correct the condition that caused the abort and retry the failing command.
- Sense Key = B (Aborted Command) with Additional Sense Code of 44 or 48 are drive caused abort conditions. The initiator should:
 1. Retry the failing command.
 2. If the retry of the failing command completes with
 - *Good Status*, log the error as recovered.
 - Abort Command Sense, verify there is no outside cause (e.g. power supply) for the failure.
 3. Retry the failing command.
 4. If the retry of the failing command completes with
 - *Good Status*, log the error as recovered.
 - Abort command sense, then service the drive using the service guideline recommended in Section 19.21.1, “Drive Service Strategy” on page 301.
- Sense Key = B (Aborted Command) and an Additional Sense Code of 47 can be an initiator or Drive caused abort condition. The initiator should follow the above procedure for initiator caused abort conditions if the Drive detected the SCSI bus parity error. The initiator should follow the above procedure for Drive caused abort conditions if the initiator detected the SCSI bus parity error.

19.21.4.8 Unit Attention Condition

Unit Attention Conditions are not errors. They alert the initiator that the drive had an action that may have changed an initiator controlled state in the drive. These conditions are the following:

Not Ready to Ready Transition

Not ready to ready transition, unit formatted. This *Unit Attention Condition* will not be reported to the initiator that issued the *Format Unit (04)*.

Reset

Reset - This means the drive was reset by either a power-on reset, LIP Reset, Target Reset or an internal reset.

Mode Parameters Changed

A *Mode Select (15)* command successfully completed. This means that the mode parameters that are the current value may have changed. The parameters may or may not have changed but the command to change the parameters successfully completed. The Drive does not actually compare the old current and the new current parameters to determine if the parameters changed. This *Unit Attention Condition* will not be reported to the initiator that issued the *Mode Select (15)*.

Microcode Has Changed

Write Buffer (3B) to download microcode has successfully completed. This means that the microcode that controls the Drive has been changed. The code may or may not be the same as the code currently being executed. The Drive does not compare old level code with new code.

Commands Cleared by Another Initiator

Tagged commands cleared by a clear queue message. This means that the command queue has been cleared. The *Unit Attention Condition* is not reported to the initiator that issued the clear queue message. *Unit Attention Condition* is reported to all initiators that had commands active or queued.

Reissue any outstanding command.

Log Select Parameters Changed

A Log Select (4C) command successfully completed. This means that the Log Select command cleared statistical information successfully (See Section 17.6, “LOG SELECT (4C)” on page 124). Unit Attention Condition is reported to all initiators excluding the initiator that issued the Log Select command.

Device Identifier Changed

A Set Device Identifier (A4) command successfully completed. This means that the Set Device Identifier information field has been updated. (See 17.41, “SET DEVICE IDENTIFIER (A4/06)” on page 237) A Unit Attention Condition is reported to all initiators excluding the initiator that issued the Set Device Identifier command.

19.21.4.9 Components Mismatch

The compatibility test is performed at a power cycle. The compatibility test verifies the microcode version of the electronics. When the Drive detects the microcode version mismatch, the most likely cause is the result of incorrect parts used during a service action.

If the error reported is Sense Key/code/qualifier 4/40/80, Diagnostic failure, bring-up fail, the initiator should do the following:

1. Retry Power cycle
2. Check the send diagnostic end status. If the status is
 - GOOD, Return to normal processing
 - *Check Condition Status*, issue a *Request Sense (03)* and follow the recommendations for the sense data returned unless the sense data is for a component mismatch. If the sense data is for component mismatch, service the drive using the service guideline recommended in Section 19.21.1, “Drive Service Strategy” on page 301.

19.21.4.10 Self Initiated Reset

The Drive will initiate a self reset when the condition of the Drive cannot be determined. The internal reset will terminate any outstanding commands, release any reserved initiators, and reset the firmware. The initiator can recover by

1. Logging the error
2. Retrying the failing command. If the failing command completes with:
 - *Good Status*, return to normal processing
 - Self initiated reset sense, service the drive according the guidelines recommended in Section 19.21.1, “Drive Service Strategy” on page 301.
 - Other, follow the recommendations for the error reported.

19.21.4.11 Defect List Recovery

This is not an error condition.

The initiator either requested a defect list in a format (block or vendor specific) that the Drive does not support or the requested defect list(s) exceed the maximum list length that can be returned. If the Sense Key/Code/Qualifier are:

1/1F/00, the requested list(s) exceed the maximum length that can be supported. The initiator should request one list at a time. If a single list exceeds the maximum returnable length, this may be an indication of a marginally operational drive. Service the drive following the service guidelines in Section 19.21.1, “Drive Service Strategy” on page 301.

1/1C/01 or 1/1C/02, the requested defect list is not in the format that the Drive supports. The requested defect list is returned in the physical format. This is the default format. There is no initiator action required for this condition.

19.21.4.12 Mismatch Recovery

A mismatch can occur on a *Verify (2F)* command or a *Write and Verify (2E)* with the byte check (ByteChk) bit active. Recovery for a mismatch error is different for the two commands.

Verify Command

The initiator should do the following:

1. Verify that the data sent to the drive is the correct data for the byte-by-byte compare.
2. Read the data from the media with a *Read (08)* or *Read (28)* command and verify that the data from the media is the expected data for the byte-by-byte compare.
 - If all data are correct, this is an indication that the data may have been read from the media incorrectly without an error detected. Service the drive using the procedure specified in Section 19.21.1, “Drive Service Strategy” on page 301.
 - If all data are not correct, this is an indication that the data on the media is not the data the initiator expected. Rewrite the correct data to the media.

Write and Verify Command

The drive uses the same data in the data buffer to write then read and compare. A mismatch error on the *Write and Verify (2E)* command is an indication that the drive cannot reliably write or read the media. Service the drive using the procedures specified in Section 19.21.1, “Drive Service Strategy” on page 301.

19.21.4.13 Microcode Error

The microcode from the interface is validated before the device operates using that microcode. When the validation detects incorrect or incomplete data, the Drive enters degraded mode.

If the initiator attempted to load microcode using the *Write Buffer (3B)* retry the *Write Buffer (3B)*. If the command completes with

- *Good Status* - return to normal processing
- *Check Condition Status* - service the drive using the service guidelines recommended in Section 19.21.1, “Drive Service Strategy” on page 301.

If the check sum error occurred during normal processing, the initiator may attempt to load microcode before deciding to service the drive using the service guidelines recommended in Section 19.21.1, “Drive Service Strategy” on page 301.

To load new microcode, the initiator should issue a *Write Buffer (3B)* command with the download and save option. If the *Write Buffer (3B)* command completes with

- Good Status, return to normal processing. Retry the failing command. If the task complete with
 - Good Status - Continue normal processing.
 - Check Condition Status for check sum error - Service the drive using the service guidelines recommended in Section 19.21.1, “Drive Service Strategy” on page 301.
 - Check Condition Status for any other error - follow the recommended recovery procedure for the error reported.
- Check Condition Status for Check sum error, service the drive using the service guidelines recommended in Section 19.21.1, “Drive Service Strategy” on page 301.
- Check Condition Status for any other error, follow the recommendations for the returned sense data.

19.21.4.14 Predictive Failure Analysis

The Drive performs error log analysis and will alert the initiator of a potential failure. The initiator should determine if this device is the only device with error activity.

If this drive is the only drive attached to the initiator with error activity, service the drive using the procedures specified in Section 19.21.1, “Drive Service Strategy” on page 301.

Note: Service for this drive can be deferred. The longer service is deferred, the more probable a failure can occur that will require immediate service.

If more than this drive is experiencing error activity, the drive is probably not at fault. Locate and service the outside source causing error activity on this drive.

20.0 SCSI Sense Data

20.1 SCSI Sense Data Format

Format of the sense data returned by the drive in response to the REQUEST SENSE command.

Table 273: Format of Sense Data.

Byte	Bit							
	7	6	5	4	3	2	1	0
0	Valid	Error Code (70h or 71h)						
1	RSVD = 0							
2	0	ILI	0	Sense Key				
3-6	(MSB)	Information Bytes						(LSB)
7	Additional Sense Length							
8-11	(MSB)	Product Specific Information						(LSB)
12	Additional Sense Code							
13	Additional Sense Code Qualifier							
14	FRU = 0							
15	SKSV	Sense-Key Specific Bits						
16-17	Sense-Key Specific Bytes							
18-19	Reserved = 0							
20-23	Vendor unique Error information							
24-29	Product Specific Information							
30-31	Reserved = 0							

20.2 Sense Data Description

20.2.1 Valid (Bit 7 of byte 0)

- 0 The Information Bytes (byte 3 through 6) are not defined.
- 1 The Information Bytes (byte 3 through 6) contain a valid logical block address.

20.2.2 Error Code (Bit 6 - 0 of byte 0)

- 70h Current Error. This indicates an error for the current command.
- 71h Deferred Error. This indicates that the error is for a previous command that has already returned a good status. Such commands are associated with the immediate bit or write caching. Format unit (04h) command is an example of a command that may return a deferred error.

20.2.3 ILI: Incorrect Length Indicator (Bit 5 of byte 2)

The ILI bit is valid for the Read Long (3Eh) command and Write Long (3Fh) command only. ILI set to one and Valid Bit set to one indicates that the requested logical block length does not match the logical block length of the data on the medium for a Read Long or Write Long command. The Information field contains residue information about the error. ILI set to zero indicates there is no incorrect length condition.

- 0 No Incorrect Length condition.
- 1 Incorrect Length Indicated.

Valid	ILI	Command = Read Long or Write Long	Description
x	0	x	No incorrect length condition
1	1	yes	Requested Logical block Length does not match the logical block length of the data on the drive

20.2.4 Sense Key (Bit 3 - 0 of byte 2)

The sense key provides generic categories in which error and exception conditions can be reported. Initiators would typically use sense keys for high level error recovery procedures.

0h	No Sense	There is no sense key information to be reported for the logical unit.
1h	Recovered Error	The last command completed successfully with some recovery action performed by the drive. More detailed information is available in the Additional Sense Code and Additional Sense Code Qualifier.
2h	Not Ready	The logical unit addressed cannot be addressed. More detailed information is available in the Additional Sense Code and Additional Sense Code Qualifier.
3h	Medium Error	The command terminated with an unrecoverable error condition caused by a flaw in the media or an error in the recorded data. More detailed information is contained in the Additional Sense Code and Additional Sense Code Qualifier.
4h	Hardware Error	The drive detected an unrecoverable hardware error while performing a command or during a diagnostic test. More detailed information is contained in the Additional Sense Code and Additional Sense Code Qualifier.
5h	Illegal Request	There was an illegal parameter in the command descriptor block or additional parameter supplied as data. If an invalid parameter is found in the CDB, then the command is terminated without altering the medium. If an invalid parameter is found in parameters supplied as data, then the drive might have altered the medium.
6h	Unit Attention	Indicates that the drive entered in the 'Unit Attention Condition'. (See Section 19.6.4, "Unit Attention Condition" on page 281)
7h	Data Protect	
8h	Not used	
9h	Vendor Specific	
Ah	Not used	
Bh	Aborted command	The drive aborted the command.
Ch-Dh	Not Implemented	
Eh	Miscompare	
Fh	Reserved	

20.2.5 Information Bytes (Byte 3 through 6)

This field is only valid when Valid Bit is one.

- ILI = 0: This field contains the unsigned LBA associated with the sense key. The LBA reported will be within the LBA range of the command as defined in the CDB.
- ILI = 1: This field contains the difference (residue) of the requested length in bytes. Negative values are indicated by two's complement notation.

Valid	ILI	Description
0	x	0x00000000 - (not used/invalid)
1	0	LBA
1	1	Residue of the requested length in bytes

20.2.6 Additional Sense Length (Byte 7)

Indicates the remaining number of bytes in the sense data. (It is always set to 18h.)

20.2.7 Command Specific Information (Byte 8 through 11)

This field is unused and will be set to zero.

20.2.8 Additional Sense Code/Qualifier (Byte 12 and 13)

The following table shows the description of the combination of Sense Key / Sense Code / Qualifier.

Valid Sense Key, Code, Qualifier Combinations Used by the Drive.			
Key	Code	Qual	Description
			Sense Key = No Sense
0	00	00	No Additional Sense Information
			(00 00) No Error.
			Sense Key = Recovered Error
0	0B	01	Temperature Warning Error
			1A02 Temperature Warning
0	0B	03	Background Selftest Failure Warning
			1A03 Background selftest failure warning
0	0B	04	Background Pre-Scan Failure Warning
			1A04 Background pre-scan failure warning
0	0B	05	Background Media Scan Failure Warning
			1A05 Background media scan failure warning
0	0B	06	Wear Warning
			1A06 Wear Warning
1	0B	01	Temperature Warning Error
			2A02 Temperature Warning
1	0B	03	Background Selftest Failure Warning
			2A03 Background selftest failure warning
1	0B	04	Background Pre-Scan Failure Warning
			2A04 Background pre-scan failure warning
1	0B	05	Background Media Scan Failure Warning
			2A05 Background media scan failure warning
1	0B	06	Wear Warning
			2A06 Wear Warning

1	5D	10	Predictive Failure Analysis Threshold Reached
			xAFE SMART: Extreme over-temperature warning
			xAFF SMART: Test warning
1	5D	53	Remaining Reserve 1
			xA53: Remaining Reserve 1
1	5D	54	Remaining Reserve 2
			xA54: Remaining Reserve 2
1	10	01	Recovered Guard check error
			17BC Recovered Guard check error
1	10	02	Recovered Application tag error
			17BA Recovered Application tag error
1	10	03	Recovered Reference tag error
			17B8 Recovered Reference tag error
1	17	01	Recovered Data with Retries
			172A XOR Rebuild successful
			172B Media Hard Error but XOR Retry recovered
			172C Recovered media error (with ECC)
1	1F	00	Partial Defect List Transfer
			1749 Partial defect list transferred
1	44	00	Internal Target Failure
			F123 Invalid request to enter sleep mode
			F128 DRAM test in progress
			F129 DRAM test complete
			1201 Error in UEC class
			1202 Error in UEC cause
			171C Recovered DRAM CRC error
			F427 HMI: DRD HMT timeout
			17C2 Recovered Read CRC error
			17C4 Recovered DRAM ECC error

			17C6 Recovered DRAM ECC LBA error
1	0B	01	Temperature Warning Error
			2A02 Temperature Warning
1	0B	03	Background Selftest Failure Warning
			2A03 Background selftest failure warning
1	0B	04	Background Pre-Scan Failure Warning
			2A04 Background pre-scan failure warning
1	0B	05	Background Media Scan Failure Warning
			2A05 Background media scan failure warning
1	0B	06	Wear Warning
			2A06 Wear Warning
			Sense Key = Not Ready
2	04	00	Logical Unit Not Ready
			F501 Logical unit not ready
2	04	01	Logical Unit Is In The Process of Becoming Ready
			F502 Logical unit becoming ready
2	04	02	Logical Unit Not Ready, initializing command required
			F124 Bring-up error
			F503 Logical unit not ready - initializing command required
2	04	03	Logical Unit Not Ready, Manual Intervention Required
			F572 Not ready - manual intervention required
2	04	04	Logical Unit Not Ready, Format In Progress
			F504 Not ready - format in progress
2	04	09	Not Ready - Self-test In Progress
			F505 Not ready - self-test in progress
2	04	11	Not Ready - Notify (Enable Spin-up) Required
			F553 Not ready - Notify (Enable Spin-up) required (SAS only)
2	04	F0	Vendor Unique - Logical Unit Not Ready
			F133 BATS error: Vendor ID mismatch
2	31	00	Medium Format Corrupted - Reassign Failed

			F506 Reassign failed
2	31	01	Format Command Failed
			F507 Format failed
3	10	01	Unrecovered Guard check error
			F7BD Unrecovered Guard check error
3	10	02	Unrecovered Application tag error
			F7BB Unrecovered Application tag error
3	10	03	Unrecovered Reference tag error
			F7B9 Unrecovered Reference tag error
3	11	00	Unrecovered Read Error
			F421 HMI: DRD CRC error
			F52D Buffer CRC error on read
			F72D Unrecovered media error
3	11	14	Unrecovered LBA write uncorrectable error
			F7CC: Unrecovered LBA write uncorrectable error
3	31	00	Medium Format Corrupted Reassign Failed
			F701 Format corrupted
4	35	00	Enclosure Services Failure
			F539 ESI: unspecified failure (FC-AL only)
4	35	01	Enclosure Services Failure - Unsupported Enclosure Function
			F53A ESI: unsupported function (FC-AL only)
4	35	02	Enclosure Services Failure - Enclosure Services Unavailable
			F53B ESI: enclosure unavailable (FC-AL only)
4	35	03	Enclosure Services Failure - Enclosure Services Transfer Failure
			F556 ESI: transfer failed - write ack
			F557 ESI: transfer failed - read ack
			F558 ESI: transfer failed - write ready
			F559 ESI: transfer failed - read ready
			F55E ESI: transfer failed - EDV
			F579 ESI: Transfer Checksum Error

			F57A ESI: Transfer Checksum Ready Timeout
4	35	04	Enclosure Services Failure - Enclosure Services Refused
			F55A ESI: transfer refused - write ack
			F55B ESI: transfer refused - read ack
			F55C ESI: transfer refused - write ready
			F55D ESI: transfer refused - read ready
4	35	80	Enclosure Services Failure
			F53C ESI: SES Timeout
4	3E	03	Self-test Failed
			F75D Self-test failed
4	3E	04	Media Self-test Failed
			F762 Media selftest hard cache fail
			F763 Media selftest OTF cache fail
4	40	80	Diagnostic Failure
			F101 BATS error: Reserved Area - Invalid request
			F102 BATS error: Reserved Area - Broken
			F103 BATS error: Reserved Area - Invalid version
			F104 BATS error: Reserved Area - Invalid checksum
			F105 BATS error: Reserved Area - Invalid eyecatcher
			F106 BATS error: Reserved Area - Invalid main header checksum
			F107 BATS error: Reserved Area - Invalid read length
			F108 BATS error: Reserved Area - Address boundary error
			F10E BATS error: Directory broken
			F10F BATS error: Overlay code load error
			F110 BATS error: Overlay code check
			F111 BATS error: RAM code load error
			F112 BATS error: RAM code check
			F113 BATS error: Config invalid
			F114 BATS error: Log manager invalid

			F11D Incorrect drive Code
			F121 BATS error: Code download in progress
			F122 BATS error: Performance data read error
			F125 BATS error: Invalid RID/FID
			F12B BATS error: Reserved area - invalid model
			F12C BATS error: Invalid code size
			F12E Format Reserved: Insufficient DIRS good
			F12F Format Reserved: Insufficient FATS good
			F131 Flash timeout
			F137 Flash ECC error
			F139 Format Reserved: Resize RIDFID error
			F13B SW target broken
			F13C NCDE DRAM Error
			F13C NCDE DRAM Error
4	40	81	DRAM Failure
			F12A DRAM test error
4	40	A0	Diagnostic Failure
			F11B BATS#2 error: CRC test error
			F11C BATS#2 error: XOR test error
			F136 BATS#2 error: End-to-End test error
4	44	00	Internal Target Failure
			F140 Cap charge time exceeded
			F141 Cap discharge time exceeded
			F142 Cap discharge time too short
			F143 Cap charge exit check failed
			F208 Nand missing
			F209 Sanity blocksize error
			F20A Mode page structure mismatch
			F60E Fconfig token update failed after code download

			F401 HMI: HMT error invalid
			F402 HMI: AES side slots overflow
			F403 HMI: HMT side valids overflow
			F404 HMI: FIFO num valid overflow
			F405 HMI: FIFO num available overflow
			F601 South: Boot incomplete
			F602 South: Trapped
			F603 South: Timeout
			F604 South: Command done
			F605 South: Command error
			F606 South: Unknown event
			F607 South: Generic assert
			F608 South: Identify failed
			F609 South: Assert dump invalid
			F60A South: Assert collision
			F60B South: Not ready for asserts
			F60C South: Dump erase but no assert found
			F60D South: Dump read but no assert found
			F60F Fconfig token parsing failed
			F610 South: Nand unsupported
			F611 South: NandID mismatch
			F612 South: No firmware found
			F613 South: Bad firmware checksum
			F620 South: Enable logical
			F621 South: Disable logical no context
			F622 South: Disable logical bad context
			F623 South: Disable logical asserted
			F624 South: Disable logical no defect map
			F625 South: Disable logical no space

			F626 South: Disable logical channel CE conflict
			F627 South: Read only
			F628 South: Log Invalid
			F71D Unrecovered DRAM CRC error
			F7C3 Unrecovered Read CRC error
			F7C5 Unrecovered DRAM ECC error
			F7C7 Unrecovered DRAM ECC LBA error
			FCxx Unable to read RID or FID number xx
4	81	00	Vendor Unique - Internal Logic Error
			F56F Log dump data memory error
5	15	00	Phy Test Function in Progress
			F50D Phy test function in progress
5	1A	00	Parameter List Length Error
			F820 Parameter list length error
5	20	00	Invalid Command Operation Code
			F821 Invalid opcode in CDB
5	21	00	Logical Block Address out of Range
			F822 LBA out of range
5	24	00	Invalid Field in CDB
			F823 Illegal request - invalid field in CDB
5	25	00	Logical Unit Not Supported
			FA24 Invalid LUN
5	26	00	Invalid Field in Parameter List
			F825 Illegal request - invalid field in parameter list
			F826 Unsupported log page
5	26	02	Parameter Value Invalid
			F120 BATS error: Code compatibility failure
			F126 BATS error: Code checksum error
			F127 BATS error: Invalid header

			F130 Incorrect Customer code
5	26	04	Invalid Release of Active Persistent Reservation
			F828 Invalid release of persistent reservation
5	2C	00	Illegal Request Sequence error
			F511 Illegal Request Sequence error
5	49	00	Invalid Message Error
			F512 Invalid message
5	55	04	Insufficient Registration Resources
			F567 Insufficient registration resources
6	0B	01	Temperature Warning Error
			3A02 Temperature Warning
6	0B	03	Background Selftest Failure Warning
			3A03 Background selftest failure warning
6	0B	04	Background Pre-Scan Failure Warning
			3A04 Background pre-scan failure warning
6	0B	05	Background Media Scan Failure Warning
			3A05 Background media scan failure warning
6	0B	06	Wear Warning
			3A06 Wear Warning
6	28	00	Not Ready To Ready Transition (Format completed)
			F514 Not ready to read transition
6	29	00	Unit Attention - Login Reset
			F515 Login reset (FC-AL only)
6	29	01	Unit Attention - POR Occurred
			F516 Power on reset
6	29	02	Unit Attention - SCSI Bus Reset Occurred
			F517 LIP Reset (FC-AL), SAS Hard Reset (SAS)
6	29	03	Unit Attention - Bus Device Reset Occurred
			F518 Target Reset (FC-AL), LUN Reset (SAS)
6	29	04	Unit Attention - Self Initiated Reset Occurred

			F519 Self initiated reset
6	29	05	Transceiver Changed to SE
			F548 Device Control Hard Reset received
6	29	07	I_T Nexus Loss Occurred
			F554 I_T Nexus Loss Occurred (SAS only)
6	2A	01	Mode Parameters Changed
			F51C Mode parameters changed
6	2A	02	Log Parameters Changed
			F51D Log parameters changed
6	2A	03	Reservations Preempted
			F51E Reservations pre-empted
6	2A	04	Reservations Released
			F51F Reservations released
6	2A	05	Registrations Released
			F520 Registrations pre-empted
6	2F	00	Commands Cleared by Another Initiator
			F521 Commands cleared by another initiator
6	2F	01	Commands Cleared by Power Loss Notification
			F573 Commands cleared due to power failure event (SAS)
6	3F	01	Microcode has been changed
			F522 Microcode changed
6	3F	03	Inquiry Parameters Changed
			F523 Inquiry parameters changed
6	3F	05	Device Identifier Changed
			F537 Device identifier changed
B	0E	02	Information Unit Too Long
			F562 Information unit too long.
B	10	00	Aborted Command CRC error
			F417 HMI: DWT CRC error
			F52C Drive CRC error

B	10	01	Aborted Command – End-to-End Guard Check
			F414 HMI: HWT guard check error
			F424 HMI: DRD guard check error
			F529 Drive Guard check error
			F568 End-to-End Data Protection Guard check
B	10	02	Aborted Command – End-to-End Application Tag Check
			F415 HMI: HWT app check error
			F423 HMI: DRD application tag check error
			F52A Drive application tag check error
			F569 End-to-End Data Protection Application Tag check
B	10	03	Aborted Command – End-to-End Reference Tag Check
			F416 HMI: HWT ref check error
			F422 HMI: DRD ref check error
			F52B Drive reference tag check error
			F56A End-to-End Data Reference Tag check
B	3F	0F	Aborted Command - Echo Buffer Overwritten
			F544 Echo buffer overwritten
B	44	00	Internal Target Failure
			F406 HMI: Abort
			F410 HMI: Write SG abort
			F411 HMI: DWT DRAM ECC error
			F412 HMI: HWT sync CRC error
			F413 HMI: HWT sync CRC LBA error
			F418 HMI: DWT HMT error
			F419 HMI: DWT HMT timeout
			F420 HMI: Read SG abort
			F425 HMI: DRD HMI LBA error
			F427 HMI: DRD HMT timeout
			F426 HMI: DRD HMT error

			F526 Drive HMI LBA error
			F527 Drive HMI error
			F528 Drive HMI Timeout
			F52E Internal target failure - Host Interface
			F54A Xfer Ready credit exceeded (FC-AL only)
			F54B Xfer length error (FC-AL only)
			F56B ECC error in DRAM customer data area
			F56C Uncorrectable DRAM ECC error
			F56E Log dump south error
			F570 Host interface CRC error
			F741 Media overall command timeout not dispatched
			F742 Media overall command timeout in recovery
			F743 Media overall command timeout executing
			F57B Host Interface Synchronous CRC LBA Error
			F75C Internal media access timeout
B	47	01	Data Phase CRC Error
			F54E Data Phase CRC Error
B	4B	00	Data Phase Error
			F53E Data phase error
B	4B	01	Invalid Target Port Transfer Tag Received
			F561 Information unit too short (SAS only)
B	4B	02	Too Much Write Data
			F560 Too much write data (SAS only)
B	4B	03	ACK/NAK Timeout
			F551 ACK/NAK Timeout (SAS only) F57D Break Received (SAS only)
B	4B	04	NAK Received
			F550 NAK Received (SAS only)
B	4B	05	Data Offset Error
			F552 Bad parameter offset (SAS only)

B	4B	06	Initiator Response Timeout
			F555 Initiator response timeout (SAS only)
B	4E	00	Overlapped Commands Attempted
			F534 Overlapped command attempted
B	4F	00	Command Aborted Due To Loop Initialization
			F53F Abort by LIP (FC-AL only), Abort by OOB (SAS)
			Sense Key = Miscompare
E	1D	00	Miscompare During Verify Operation
			F535 Miscompare during verify

20.2.9 RU: Field Replaceable Unit (Byte 14)

The FRU (Field Replaceable Unit) field value will always be zero.

20.2.10 Sense Key Specific (Byte 15 through 17)

The definition of this field is determined by the value of the sense key field.

20.2.10.1 Sense Key Specific - Illegal Request (Sense Key = 5h)

Error field pointer is returned.

Table 274: Field Pointer Bytes

Byte	Bit							
	7	6	5	4	3	2	1	0
15	SKSV	C/D	Reserved		BPV	Bit Pointer		
16 17	(MSB) Field Pointer (LSB)							

SKSV Sense-key specific valid

0 Sense-key specific field is not valid.

1 Sense-key specific field is valid.

C/D Command/Data

0 Indicates that the illegal parameter was in the data parameters sent by the initiator during DATA OUT phase

1 Indicates that the illegal parameter was in the command descriptor block.

BPV Bit Pointer Valid

0 Bit pointer field is not valid.

1 Bit pointer field is significant.

Bit Pointer Indicates which bit of the byte number reported in Field Pointer is the bit in error. When a multiple bit field is in error, the pointer points to the most significant bit of the field.

Field Pointer Indicates which bytes of the command descriptor block or of the parameter data were in error. Bytes are numbered starting from zero, as shown in the tables describing the commands and parameters. When a multiple byte field id is in error, the pointer points to the most significant byte of that field.

20.2.10.2 Sense Key Specific - Recovered (Sense Key = 1h) or Medium (Sense Key = 3h) or Hardware (Sense Key = 4h)

Hardware (Sense Key = 4h) or Medium Error (Sense Key = 3h)

Actual Retry Count is reported.

Table 275: Actual Retry Count

Byte	Bit							
	7	6	5	4	3	2	1	0
15	SKSV	Reserved						
16	Reserved							
17	Actual Retry Count							

SKSV

Sense-key specific valid

0 Actual Retry Count is not valid.

1 Actual Retry Count is valid.

Actual Retry Count

Number of retry steps used in attempting to recover from the error condition.

20.2.10.3 Sense Key Specific - Not Ready (Sense key = 2h)

These fields are defined for the Format unit (04h) command with the Immediate bit set to one and the Send Diagnostic (1Dh) command with Background self-test function.

Progress indication is returned.

Table 276: Progress Indication

Byte	Bit							
	7	6	5	4	3	2	1	0
15	SKSV	Reserved						
16	(MSB) Progress Indication (LSB)							
17								

SKSV Sense-key specific valid

0 Progress Indication is not valid.

1 Progress Indication is valid.

Progress Indication Indicates a percent complete in which the returned value is the numerator that has 10000h as its denominator.

20.2.11 Reserved (Byte 18 through 19)

Reserved fields are filled with zero.

20.2.12 Vendor unique error information (Byte 20 through 23)

This field gives detailed information about the error. It contains a unique code which describes where the error was detected and which piece of hardware or microcode detected the error depending on current operation.

20.2.13 Physical Error Record (Byte 24 thru 29)

- ILI = 1 - This field contains zeros.
- ILI = 0 - These bytes contain the physical location of the error.

If physical location has no relevance for the error, bytes 24 through 29 will all be set to 0FFFFFFFFFh for Valid = 0 and ILI = 0. This Physical Error Record field is valid for Sense Key 1, 3, and 4 only.

Valid	ILI	Description
1	0	Die (bytes 24-25) Channel (byte 26) CE (byte 27) Sector Offset (byte 28) Page (byte 29)
1	1	0x000000000000
0	x	0x000000000000 - (not used/invalid)

20.2.14 Reserved (Byte 30 through 31)

Reserved fields are filled with zero.

21.0 Appendix. UEC list

Following is the list of Unit Error Codes and associated descriptions. The Unit Error Codes are returned by the target in sense data bytes 20-21.

The list of Unit Error Codes and descriptions does not have a direct correlation to the error descriptions and Sense Key/Code/Qualifier descriptions in Section 19.0, “Additional information” on page 273. These codes are used internally by HGST and may change without notice.

How to find a specific UEC

The second hex digit indicates the grouping, e.g. interface, media, servo, etc. types of errors. The table is sorted without regard to the first hex digit; instead, sorting is by the least significant three hex digits.

Table 277: Unit Error Codes

UEC	Description
0000	No error
F101	BATS error: Reserved Area - Invalid request
F102	BATS error: Reserved Area - Broken
F103	BATS error: Reserved Area - Invalid version
F104	BATS error: Reserved Area - Invalid checksum
F105	BATS error: Reserved Area - Invalid eyecatcher
F106	BATS error: Reserved Area - Invalid main header checksum
F107	BATS error: Reserved Area - Invalid read length
F108	BATS error: Reserved Area - Address boundary error
F10E	BATS error: Directory broken
F10F	BATS error: Overlay code load error
F110	BATS error: Overlay code check
F111	BATS error: RAM code load error
F112	BATS error: RAM code check
F113	BATS error: Config invalid
F114	BATS error: Log manager invalid
F119	Bats#2 error: Read Write test error
F11B	BATS#2 error: CRC test error
F11C	BATS#2 error: XOR test error
F11D	Incorrect drive Code
F120	BATS error: Code Compatibility Failure
F121	BATS error: Code download in progress
F122	BATS error: Performance data read error
F123	Invalid request to enter sleep mode
F124	Bring-up error
F125	BATS error: Invalid RID/FID
F126	BATS error: Code checksum error
F127	BATS error: Invalid header
F128	DRAM test in progress

UEC	Description
F129	DRAM test complete
F12A	DRAM test error
F12B	BATS error: Reserved area - invalid model
F12C	BATS error: Invalid code size
F12E	Format Reserved: Insufficient DIRS good
F12F	Format Reserved: Insufficient FATS good
F130	Incorrect Customer Code
F131	Flash time-out
F133	BATS error: Vendor ID mismatch
F136	BATS#2 error: End-To-End Data Protection error
F137	Flash ECC error
F139	Format Reserved: Resize RIDFID error
F13B	SW Target broken
F13C	NCDE DRAM failure
F140	BATS error: Cap charge time exceeded
F141	BATS error: Cap discharge time exceeded
F142	BATS error: Cap discharge time too short
F143	BATS error: Cap charge exit check failed
1201	Error in UEC class
1202	Error in UEC cause
F203	Internal target failure
F207	AHB Access Error
F208	NAND Missing
F209	North/South disagree on block size
F20A	Mode page structure mismatch
F401	HMI: HMT error invalid
F402	HMI: AES side slots overflow
F403	HMI: HMT side valids overflow
F404	HMI: FIFO num valid overflow
F405	HMI: FIFO num available overflow
F406	HMI: Abort
F410	HMI: Write SG abort
F411	HMI: DWT DRAM ECC error
F412	HMI: HWT sync CRC error
F413	HMI: HWT sync CRC LBA error
F414	HMI: HWT guard check error
F415	HMI: HWT app check error
F416	HMI: HWT ref check error

UEC	Description
F417	HMI: DWT CRC error
F418	HMI: DWT HMT error
F419	HMI: DWT HMT time-out
F420	HMI: Read SG abort
F421	HMI: DRD CRC error
F422	HMI: DRD ref check error
F423	HMI: DRD app check error
F424	HMI: DRD guard check error
F425	HMI: DRD HMI LBA error
F426	HMI: DRD HMT error
F427	HMI: DRD HMT time-out
F501	Logical unit not ready
F502	Logical unit becoming ready
F503	Logical unit not ready - initializing command required
F504	Not ready - format in progress
F505	Not ready - self-test in progress
F506	Reassign failed
F507	Format failed
F50D	Phy Test Function in Progress
F511	Illegal request sequence error
F512	Invalid Message
F514	Not ready to ready transition
F515	Login reset (FC-AL only)
F516	Power on reset
F517	LIP reset (FC-AL), SAS Hard Reset (SAS)
F518	Target Reset (FC-AL), LUN Reset (SAS)
F519	Self initiated reset
F51C	Mode parameters changed
F51D	Log parameters changed
F51E	Reservations pre-empted
F51F	Reservations released
F520	Registrations pre-empted
F521	Commands cleared by another initiator
F522	Microcode changed
F523	Inquiry parameters changed
F526	Drive HMI LBA error
F527	Drive HMT error
F528	Drive HMT time-out
F529	Drive guard check error
F52A	Drive application tag error

UEC	Description
F52B	Drive reference tag error
F52C	Drive CRC error
F52D	Buffer CRC error on read
F52E	Internal target failure
F534	Overlapped command attempted
F535	Miscompare during verify
F536	Reservation conflict
F537	Device identifier changed
F539	ESI: unspecified failure (FC-AL only)
F53A	ESI: unsupported function (FC-AL only)
F53B	ESI: enclosure unavailable (FC-AL only)
F53C	ESI: transfer failure (FC-AL only)
F53E	Data phase error
F53F	Abort by LIP (FC-AL), Abort by OOB (SAS)
F544	Echo buffer overwritten
F548	Device Control Hard Reset received
F54A	Xfer Ready credit exceeded (FC-AL only)
F54B	Transfer length error (FC-AL only)
F54E	Data Phase CRC Error
F550	NAK revd (SAS)
F551	ACK NAK Time-out (SAS)
F552	Bad parameter offset (SAS)
F553	LUN Not ready, Notify (Enable Spinup) required (SAS)
F554	I_T_Nexus Loss Occurred (SAS)
F555	Initiator Response Time-out (SAS)
F556	ESI transfer failed - write ack (FC-AL)
F557	ESI transfer failed - read ack (FC-AL)
F558	ESI transfer failed - write ready (FC-AL)
F559	ESI transfer failed - read ready (FC-AL)
F55A	ESI transfer refused - write ack (FC-AL)
F55B	ESI transfer refused - read ack (FC-AL)
F55C	ESI transfer refused - write ready (FC-AL)
F55D	ESI transfer refused - read ready (FC-AL)
F55E	ESI transfer failed - EDV (FC-AL)
F560	Too much write data (SAS)
F561	Information unit too short (SAS)
F562	Information unit too long (SAS)
F566	Not ready in HGST DST
F567	Insufficient registration resources
F568	End-to-End Data Protection Guard check
F569	End-to-End Data Protection Application Tag check

UEC	Description
F56A	End-to-End Data Protection Reference Tag check
F56B	ECC error in DRAM customer data area
F56C	Uncorrectable DRAM ECC error
F56E	Log dump south error
F56F	Log dump data memory error
F570	Host interface Synchronous CRC error
F572	LUN not ready; manual intervention required
F573	Commands cleared due to power failure event (SAS)
F574	Unsupported hardware
F579	ESI Transfer Checksum Error (FC-AL)
F57A	ESI Transfer Checksum Ready Time-out (FC-AL)
F57B	Host Interface Synchronous CRC LBA Error
F57D	BREAK received
F601	South: Boot incomplete
F602	South: Trapped
F603	South: Time-out
F604	South: Command done
F605	South: Command error
F606	South: Unknown event
F607	South: Generic assert
F608	South: Identify failed
F609	South: Assert dump invalid
F60A	South: Assert collision
F60B	South: Not ready for asserts
F60C	South: Dump erase but no assert found
F60D	South: Dump read but no assert found
F60E	South: Fconfig token update failed after code download
F60F	South: Fconfig token parsing failed
F610	South: Nand unsupported
F611	South: NandID mismatch
F612	South: No firmware found
F613	South: Bad firmware checksum
F614	South: No slots found
F620	South: Enable logical
F621	South: Disable logical no context
F622	South: Disable logical bad context
F623	South: Disable logical asserted
F624	South: Disable logical no defect map
F625	South: Disable logical no space
F626	South: Disable logical channel CE conflict

UEC	Description
F627	South: Read only
F628	South: Log Invalid
F701	Format corrupted
171C	Recovered DRAM CRC error
F71D	Unrecovered DRAM CRC error
172A	XOR Rebuild successful
172B	Media Hard Error but XOR Retry recovered
172C	Recovered media error (with ECC)
F72D	Unrecovered media error
F741	Media overall command time-out not dispatched
F742	Media overall command time-out in recovery
F743	Media overall command time-out executing
1749	Partial defect list transferred
F75C	Internal media access time-out
F75D	Selftest failed
F762	Cache test fail
F763	OTF cache fail
17B8	Recovered Reference tag error
F7B9	Unrecovered Reference tag error
17BA	Recovered Application tag error
F7BB	Unrecovered Application tag error
17BC	Recovered Guard check error
F7BD	Unrecovered Guard check error
17C2	Recovered Read CRC error
F7C3	Unrecovered Read CRC error
17C4	Recovered DRAM ECC error
F7C5	Unrecovered DRAM ECC error
17C6	Recovered DRAM ECC LBA error
F7C7	Unrecovered DRAM ECC LBA error
F7CC	Unrecovered LBA ECC write uncorrectable
F820	Parameter list length error
F821	Invalid opcode in CDB
F822	LBA out of range
F823	Illegal request - invalid field in CDB

UEC	Description
F824	Invalid LUN
F825	Illegal request - invalid field in parameter list
F826	Illegal request - Unsupported Log Page
F828	Illegal request - Invalid Release of Persistent Reservation
1A02	SMART: Temperature warning (no sense)
2A02	SMART: Temperature warning (recovered sense)
3A02	SMART: Temperature warning (unit attn sense)
1A03	SMART: Background selftest warning (no sense)
2A03	SMART: Background selftest warning (recovered sense)
3A03	SMART: Background selftest warning (unit attn sense)
1A04	SMART: Background Pre-Scan warning (no sense)
2A04	SMART: Background Pre-Scan warning (recovered sense)
3A04	SMART: Background Pre-Scan warning (unit attn sense)
1A05	SMART: Background Media Scan warning (no sense)
2A05	SMART: Background Media Scan warning (recovered sense)
3A05	SMART: Background Media Scan warning (unit attn sense)
1A06	SMART: Wear warning (no sense)
2A06	SMART: Wear warning (recovered sense)
3A06	SMART: Wear warning (unit attn sense)
1A28	SMART: Capacitor fail (no sense)
2A28	SMART: Capacitor fail (recovered sense)
3A28	SMART: Capacitor fail (unit attn sense)
1A53	SMART: Remaining Reserve 1 (no sense)
2A53	SMART: Remaining Reserve 1 (recovered sense)
3A53	SMART: Remaining Reserve 1 (unit attn sense)
1A54	SMART: Remaining Reserve 2 (no sense)
2A54	SMART: Remaining Reserve 2 (recovered sense)
3A54	SMART: Remaining Reserve 2 (unit attn sense)
1AFE	SMART: Thermal Sense trip (no sense)
2AFE	SMART: Thermal Sense trip (recovered sense)
3AFE	SMART: Thermal Sense trip (unit attn sense)
1AFF	SMART: Test warning (no sense)
2AFF	SMART: Test warning (recovered sense)
3AFF	SMART: Test warning (unit attn sense)
FCxx	Unable to read RID or FID number xx

A

- Abbreviations 107
- abort sequence 47
- abort task 101
- Abort Task (Implemented as ABTS BLS) 101
- abort task set 101
- Additional information 273
- Additional Sense Length (Byte 7) 312
- Address Discovery 67
- address discovery (ADISC) 67
- Alternate Sectors per Zone 163
- Alternate Tracks per Logical Unit 163
- Alternate Tracks per Zone 163
- Appendix 331
- auto-contingent allegiance (ACA) 102
- Automatic Read Reallocation Enabled 161
- Automatic Rewrite/Reallocate 292
- Automatic Write Reallocation Enabled 161
- autosense 103

B

- basic accept (BA_ACC) 48
- Basic Link Services 47
- basic link services 47
- basic reject (BA_RJT) 49
- Block Descriptor 157
- Byte ordering conventions 107

C

- capability entry(s) 77
- Caution 1
- class service parameters 58, 65
- clear ACA 102
- clear task set 102
- Command Processing During Execution of Active I/O process 279
- Command Processing During Startup and Format Operations 283
- Command processing while reserved 290
- Command Queuing 291
- Command queuing 291
- Command Reordering 291
- Command reordering 291
- Command Time out Limits 299
- Common Fibre Channel Services 85
- common service parameters 55
- common transport (CT) interface 85
- Concurrent I/O Process 291
- concurrent sequences 60
- Control Electronics 5

Cooling requirements 21
Corrosion test 20
cyclic redundancy check (CRC) 46
D
Data Integrity 11
Data Reliability 27
Data transfer speed 10
DC Power Requirements 23
defect descriptor 110
defect list 109
Deferred Error Condition 283
Degraded Mode 284
Diagnostics 295
Disable Block Descriptor 154
document identifier 78
Download Microcode and Save (Mode 0101b) 263
Drive Characteristics 7
Drive ready time 9
E
Electrical Interface 13
Electromagnetic Compatibility 37
enclosure services 293
end-to-end credit 60
Environment 19
Equipment errors 27
Equipment Status 11
European Standards Compliance 41
exchange 47
Extended Link Services 50
extended link services 50
F
Fabric 62
fabric address notification (FAN) 66
Fabric Login 62
fabric login (FLOGI) 62
Failure prediction (PFA/S.M.A.R.T) 27
FC-AL attachment 43
FC-AL Connector 13
FC-AL timers 90
FCP_CDB 96
FCP_CMND IU 95
FCP_CNTL 95
FCP_DL 96
FCP_LUN 95
FCP_RESID 98
FCP_RSP_INFO 99

FCP_RSP_LEN 98
FCP_SNS_INFO 99
FCP_SNS_LEN 98
FCP_XFER_RDY 97
Flammability 41
Flow Control 275
FORMAT UNIT 108
format unit 108
Formatted Capacity 7
frame delimiter 45
frames 46
Fundamentals 43
G
German Safety Mark 41
Glossary 1
Grown Defect List (Glist) 203
I
Identification 35
Idle Time Function 299
IEEE registered name format 44
Incorrect Length Indicator (Bit 5 of byte 2) 310
information units (IUs) 93
initiator control (IC) 59
inquiry 112
Interface Connector 30, 32
Internal Error Condition 283
invalid transmission word 75
L
Labels 35
link failure 91
link failure count 75
link service accept (LS_ACC) 51
link service reject (LS_RJT) 52
Log Page Parameters 128
LOG SELECT 124
LOG SELECT (4C) 124
LOG SENSE 127
LOG SENSE (4D) 127
Logical Block Cache Segment Size 167
Logout (LOGO) 61
logout (LOGO) 61
loop initialization frame 51
Loop Initialization Procedure 273
loop initialization procedure 273
loss of signal count 75
loss of sync count 75

M

- Maximum Burst Size 162
- Maximum Pre-fetch 167
- Maximum Pre-fetch Ceiling 167
- Mechanical Specifications 29
- Minimum Pre-fetch 167
- MODE SELECT 152, 153
- MODE SELECT (15) 152
- MODE SELECT (55) 153
- MODE SENSE (1A) 154
- MODE SENSE (5A) 178
- Multiple Initiator Environment 294
- Multiple Initiator Systems 293
- multiple initiator systems 293

N

- N_Port ID 68
- NL_Port address 44
- node name 68
- Non-operating shock 34
- Non-operating Vibrations 33

O

- Obtaining an AL_PA 273
- obtaining an AL_PA 273
- open sequences per exchange 61
- Operating shock 33
- Operating Vibration 33
- ordered set 45
- originator process associator 69
- Outline 29
- Outline of the Drive 3

P

- Packaging 20
- Page Code Valid 212
- Page Format 232
- Parameter data for Read Keys 180
- Parameter data for Read Reservations 181
- Parameter list 184
- payload length 76
- Persistent Reservation 179
- PERSISTENT RESERVE IN (5E) 179
- PERSISTENT RESERVE OUT 182, 186
- PERSISTENT RESERVE OUT (5F) 182, 186
- Persistent Reserve Out Parameter list 184
- Physical Sector Format (101b) 205
- port discovery (PDISC) 67
- port login (PLOGI) 53

port name 68
Post Error 161
Predictive Failure Analysis 308
PRE-FETCH 187
Preventive Maintenance 27
Primary Defect List 203
primitive sequence 45
primitive sequence protocol error 75
primitive signal 45
Priority Commands 290
Priority commands 290
process login (PRLI) 68
process logout (PRLO) 72
Product ID 8
public loop operation 276
Q
Queue Depth 291
Queue Full Status 291
R
Random Vibration 33
READ (10) - (28) 189
Read Ahead 293
READ BUFFER (3C) 195
Read Buffer Full Ratio 162
Read Cache Disable 167
READ CAPACITY 199
READ CAPACITY command 199
Read Continuous 161
READ DEFECT DATA (37) 203
READ DEFECT DATA (B7) 207
read link error status block (RLS) 74
READ LONG command 209
Read Retry Count 161
REASSIGN BLOCKS (07) 210
receive data field 60
RECEIVE DIAGNOSTICS (1C) 212
Recommended Initiator ERP 301
register FC-4 types (RFT_ID) 88
re-instate recovery qualifiers (RRQ) 79
RELEASE (17) 216
RELEASE (57) 217
Reliability 27
REPORT LUNS (A0) 226
report node capabilities (RNC) 75
REQUEST SENSE (03) 227
RESERVE (16) 228

- RESERVE (56) 229
- Reserved Area 214
- Reset 294
- Reset Actions 295
- responder process associator 69, 81
- REZERO UNIT (01) 230
- Ripple Voltage 25
- RNC flags 76
- S
- Safety 41
- SCSI Command Set 105
- SCSI Control Byte 107
- SCSI Protocol 279
- SCSI Sense Data 309
- SCSI Sense Data Format 309
- SCSI Status Byte 271
- SCSI Status Byte Reporting 279
- SCSI-FCP 93
- Sectors per Track 164
- SEEK (10) - (2B) 231
- SEEK (6) - (0B) 231
- Segmented Caching 293
- SEND DIAGNOSTIC (1D) 232
- sense data 293
- Sense Data Description 310
- Sense Key (Bit 3 - 0 of byte 2) 311
- SenseData 309
- sequence 46
- Service Action 179, 183
- Size Enable 167
- START STOP UNIT (1B) 238
- Storage requirements 20
- Storage time 20
- Summary 185
- Swept Sine Vibration 33
- SYNCHRONIZE CACHE (35) 239
- T
- target reset 102
- Task Attributes 99
- Task Management functions 100
- task management functions 93
- Temperature and humidity 19
- Temperature Threshold 159
- terminate task 102
- Terminology 93
- TEST UNIT READY (00) 241, 242

third party N_Port ID 81
third party originator process associator 81
third party process logout (TPRLO) 80
Track Skew Factor 164
Tracks per Zone 164
Transfer Block 161
transmission character 45
Type 183
U
UEC list 331
UL and CSA Standard Conformity 41
Unit Attention Condition 305
unit attention condition 281
V
vendor identifier 76
vendor unique information 76
Verify 307
VERIFY (2F) 244
Vibration and Shock 33
VU information length (VU_Len) 76
W
word 45
WRITE (10) - (2A) 251
WRITE (6) - (0A) 250
WRITE AND VERIFY (2E) 257, 259
WRITE BUFFER (3B) 263
Write Buffer Empty Ratio 162
Write Cache 291
Write Cache Enable 167
WRITE LONG (3F) 264
Write Retention Priority 167
Write Retry Count 161
WRITE SAME (41) 266

© Copyright HGST, a Western Digital company

HGST, a Western Digital company
3403 Yerba Buena Road
San Jose, CA 95135
Produced in the United States
Tel: 800-801-4618 Fax: 408-717-5000

05 October 2012

All rights reserved Ultrastar™ is a trademark of HGST.

HGST trademarks are authorized for use in countries and jurisdictions in which HGST has the right to use, market and advertise the brands. HGST shall not be liable to third parties for unauthorized use of HGST trademarks.

Microsoft, Windows XP, and Windows are trademarks of Microsoft Corporation in the United States, other countries, or both.

Other trademarks are the property of their respective owners.

References in this publication to HGST products, programs or services do not imply that HGST intends to make these available in all countries in which HGST operates.

Product information is provided for information purposes only and does not constitute a warranty.

Information is true as of the date of publication and is subject to change. Actual results may vary. This publication is for general guidance only. Photographs may show design models.

This publication is for general guidance only. Photographs may show design models