



OEM HARD DISK DRIVE SPECIFICATIONS

for

DALS-3540 (540 MB)

3.5-Inch Hard Disk Drive with SCSI Interface

Revision (1.3)



OEM HARD DISK DRIVE SPECIFICATIONS

for

DALS-3540 (540 MB)

3.5-Inch Hard Disk Drive with SCSI Interface

Revision (1.3)

First Edition S39H-1055-00 (March 24 1995)
Second Edition S39H-1055-01 (March 28 1995)
Third Edition S39H-1055-02 (April 13 1995)
Fourth Edition S39H-1055-03 (May 11 1995)

The following paragraph does not apply to the United Kingdom or any country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION 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. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time.

It is possible that this publication may contain reference to, or information about, IBM 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 IBM intends to announce such IBM products, programming, or services in your country.

Technical information about this product is available from your local IBM representative or at **<http://www.ibm.com/harddrive>**

IBM 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. You can send license inquiries, in writing, to the IBM Director of Commercial Relations, IBM Corporation, Armonk, NY 10577.

© Copyright International Business Machines Corporation 1995. All rights reserved.

Note to U.S. Government Users —Documentation related to restricted rights —Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract with IBM Corp.

Contents

1.0 General	1
1.1 Introduction	1
1.2 References	1
1.3 Glossary	1
1.4 General Caution	1
1.0 Outline of the drive	3
<hr/>	
Part 1. Functional Specification	5
2.0 Fixed Disk Subsystem Description	7
2.1 Control Electronics	7
2.2 Head Disk Assembly	7
2.3 Actuator	7
3.0 Drive Characteristics	9
3.1 Formatted Capacity	9
3.2 Data Sheet	9
3.3 Performance Characteristics	10
3.3.1 Command Overhead	10
3.3.2 Mechanical Positioning	11
3.3.3 Drive Ready Time	13
3.3.4 Data Transfer Speed	13
3.3.5 Buffering Operation (Lookahead/Write Cache)	13
3.3.6 Throughput	14
3.3.7 Operating Mode Definition	15
4.0 Data integrity	17
4.1 Equipment Status	17
4.2 Error Recovery	17
5.0 Physical Format	19
5.1 Shipped Format	19
6.0 Specification	21
6.1 Electrical interface specification	21
6.1.1 Connectors	21
6.1.2 Signal Definition	22
6.1.3 SCSI Cable	22
6.1.4 SCSI Bus Terminator	22
6.1.5 SCSI Bus Electrical Characteristics	23
6.1.6 The pin assignment	24
6.1.7 Signal description	25
6.1.8 Shipping Default Condition	25
6.2 Environment	26
6.2.1 Temperature and Humidity	26
6.3 DC Power Requirements	27
6.4 Reliability	28
6.4.1 Mean Time Between Failures (MTBF)	28
6.4.2 Contact Start Stop (CSS)	28

6.4.3	Useful Life	28
6.4.4	Preventive Maintenance	28
6.4.5	Data Reliability	28
6.4.6	Seek/ID Mis-compare Errors	28
6.4.7	Equipment Errors	28
6.5	Mechanical Specifications	29
6.5.1	Outline	29
6.5.2	Mechanical Dimensions	29
6.5.3	Connector Locations	31
6.5.4	Connector and Jumper	32
6.5.5	Drive Mounting	34
6.5.6	Shipping Zone and Lock	34
6.6	Vibration and Shock	35
6.6.1	Operating Vibration	35
6.6.2	Non-Operating Vibrations	35
6.6.3	Operating Shock	36
6.6.4	Non-Operating Shock	36
6.7	Acoustics	37
6.7.1	Sound Power Levels	37
6.7.2	Sound Power Acceptance Criteria	37
6.8	Identification Labels	38
6.9	Electromagnetic Compatibility	38
6.10	Safety	39
6.10.1	Underwriters Lab (UL) Approval	39
6.10.2	Canadian Standards Authority (CSA) Approval	39
6.10.3	IEC Compliance	39
6.10.4	German Safety Mark	39
6.10.5	Flammability	39
6.10.6	Safe Handling	39
6.10.7	Environment	40
6.10.8	Secondary Circuit Protection	40
6.11	Packaging	40

Part 2. SCSI Interface Specification 41

7.0	SCSI COMMAND SET	43
7.1	Abbreviations	45
7.2	Flag and Link Bits	45
7.3	Byte ordering conventions	45
7.4	FORMAT UNIT (04h)	46
7.4.1	Defect List	47
7.5	INQUIRY (12h)	49
7.5.1	INQUIRY Data Format (When Invalid LUN is Specified)	51
7.5.2	Inquiry Data Format when EVPD = 0	52
7.5.3	Inquiry Data Format when EVPD = 1 (Page Code = 00h)	55
7.5.4	Inquiry Data Format when EVPD = 1 (Page Code = 01h)	56
7.5.5	Inquiry Data Format when EVPD = 1 (Page Code = 03h)	57
7.5.6	Inquiry Data Format when EVPD = 1 (Page Code = 80h)	58
7.5.7	Inquiry Data Format when EVPD = 1 (Page Code = 82h)	59
7.6	LOG SELECT (4Ch)	61
7.7	LOG SENSE (4Dh)	61
7.8	MODE SENSE (1Ah)	62
7.8.2	Mode Parameter List	64
7.9	MODE SELECT (15h)	67

7.10	Mode select data	69
7.10.1	Page 0h (Vendor Unique Parameters)	70
7.10.2	Page 1h (Read/Write Error Recovery Parameters)	73
7.10.3	Page 2h (Disconnect/Reconnect Parameters)	78
7.10.4	Page 3h (Format Device Parameters)	79
7.10.5	Page 4h (Rigid Disk Drive Geometry Parameters)	81
7.10.6	Page 7h (Verify Error Recovery Parameters)	82
7.10.7	Page 8h (Caching Parameters)	84
7.10.8	Page Ah (Control Mode Page Parameters)	86
7.10.9	Page 0Dh (Power Condition)	88
7.11	PRE-FETCH (34h)	89
7.12	READ (08h)	91
7.13	READ CAPACITY (25h)	92
7.14	READ DEFECT DATA (37h)	94
7.14.1	Defect List Header	95
7.14.2	Bytes from Index Format (100b)	95
7.14.3	Physical Sector Format (101b)	95
7.15	READ EXTENDED (28h)	97
7.16	READ BUFFER (3Ch)	98
7.16.1	Combined Header And Data (Mode 000b)	98
7.16.2	Read Data (Mode 010b)	99
7.16.3	Descriptor (Mode 011b)	99
7.17	READ LONG (3Eh)	101
7.18	REASSIGN BLOCKS (07h)	102
7.19	RECEIVE DIAGNOSTIC (1Ch)	104
7.19.1	Receive Diagnostic Page 0	104
7.19.2	Receive Diagnostic Page 40h	104
7.20	RELEASE (17h)	106
7.21	REQUEST SENSE (03h)	107
7.22	RESERVE (16h)	108
7.22.1	Command Processing While Reserved	109
7.23	REZERO UNIT (01h)	110
7.24	SEEK (0Bh)	111
7.25	SEEK EXTENDED (2Bh)	112
7.26	SEND DIAGNOSTIC (1Dh)	113
7.26.1	Send Diagnostic Page 0	114
7.26.2	Send Diagnostic Page 40h	114
7.27	START/STOP UNIT (1Bh)	116
7.28	SYNCHRONIZE CACHE (35h)	117
7.29	TEST UNIT READY (00h)	118
7.30	VERIFY (2Fh)	119
7.31	WRITE (0Ah)	120
7.32	WRITE EXTENDED (2Ah)	121
7.33	WRITE AND VERIFY (2Eh)	122
7.34	WRITE BUFFER (3Bh)	123
7.34.1	Combined Header And Data (Mode 000b)	123
7.34.2	Write Data (Mode 010b)	124
7.34.3	Download Microcode (Mode 100b)	124
7.34.4	Download Microcode and Save (Mode 101b)	125
7.35	WRITE LONG (3Fh)	127
8.0	SCSI Status Byte	129
9.0	SCSI MESSAGE SYSTEM	131
9.1	Supported Messages	131

9.1.1	COMMAND COMPLETE (00)	132
9.1.2	SYNCHRONOUS DATA TRANSFER REQUEST (01,03,01H)	132
9.1.3	SAVE DATA POINTER (02)	134
9.1.4	RESTORE POINTERS (03)	134
9.1.5	DISCONNECT (04)	134
9.1.6	INITIATOR DETECTED ERROR (05)	134
9.1.7	ABORT (06)	135
9.1.8	MESSAGE REJECT (07)	135
9.1.9	NO OPERATION (08)	135
9.1.10	MESSAGE PARITY ERROR (09)	135
9.1.11	LINKED COMMAND COMPLETE (0A)	136
9.1.12	LINKED COMMAND COMPLETE WITH FLAG (0B)	136
9.1.13	BUS DEVICE RESET (0C)	136
9.1.14	ABORT TAG (0D)	136
9.1.15	CLEAR QUEUE TAG (0E)	136
9.1.16	QUEUE TAG MESSAGES(20h, 21h, 22h)	136
9.1.17	IDENTIFY (80 - FF)	137
9.2	Supported Message Functions	138
9.3	Attention Condition	139
9.4	SCSI Bus Related Error Handling Protocol	139
9.4.1	Unexpected BUS FREE Phase Error Condition	140
9.4.2	MESSAGE OUT Phase Parity Error	140
9.4.3	MESSAGE IN Phase Parity Error (Message Parity Error)	140
9.4.4	COMMAND Phase Parity Error	140
9.4.5	DATA OUT Phase Parity Error	141
9.4.6	INITIATOR DETECTED ERROR Message	141
9.4.7	MESSAGE REJECT Message	141
10.0	Additional Information	143
10.1	SCSI Protocol	143
10.1.1	Priority of SCSI Status Byte Reporting	143
10.1.2	Invalid LUN in Identify Message	143
10.1.3	Incorrect Initiator Connection	144
10.1.4	Command Processing During Execution of Active I/O process	144
10.1.5	Unit Attention Condition	147
10.1.6	Command Processing During Start-up and Format Operations	147
10.1.7	Internal Error Condition	148
10.1.8	Deferred error	149
10.1.9	Degraded Mode	149
10.1.10	Degraded mode handling	149
10.1.11	Command Processing While Reserved	150
10.2	Priority Commands	151
10.3	Command queuing	151
10.3.1	Queue depth	152
10.3.2	Tagged queuing	152
10.3.3	Untagged queuing	152
10.3.4	Command queuing rule	152
10.3.5	Queue Full status	153
10.3.6	Device behaviour on Command queuing	154
10.4	Command reordering	154
10.4.1	Reorder method	154
10.4.2	Command aging	155
10.4.3	The restriction of reordering	155
10.4.4	Reordering general rules	156
10.5	Concurrent I/O Process	156

10.6	Back to Back Write	157
10.7	Write Cache	157
10.8	Power Saving Mode	158
10.8.1	Overview	158
10.9	Automatic Rewrite/Reallocate	160
10.10	Segmented Caching	163
10.10.1	Overview	163
10.10.2	Read Ahead	163
10.11	Reselection	165
10.12	Single Initiator Selection	165
10.13	Non-arbitrating systems	165
10.14	Selection without ATN	166
10.15	Multiple Initiator Environment	166
10.15.1	Initiator Sense Data	166
10.15.2	Initiator Mode Select/Mode Sense Parameters	166
10.15.3	Initiator Data Transfer Mode Parameter	166
10.16	Options Jumper Block	166
10.17	Contingent allegiance Condition	167
10.18	Reset	168
10.18.1	Reset Sources	168
10.18.2	Reset Actions	168
10.19	RAS Characteristics	169
10.19.1	Power on Diagnostics	169
10.19.2	Diagnostics Command	169
10.19.3	Diagnostics Fault Reporting	170
11.0	SCSI SENSE DATA	173
11.1	SCSI Sense Data Format	173
11.1.1	Sense Data Description	174
12.0	Recovery Procedure	185
12.1	Data Recovery Procedure for Data Field Errors	185
12.1.1	Read commands	186
12.1.2	Write Commands	187
12.2	Error Recovery for Non-Data Errors	187
12.2.1	Position (Servo)	187
12.2.2	Spindle Motor Error	187
12.3	Priority of Error Reporting	188
Index		189

Figures

1.	Formatted Capacity	9
2.	Data Sheet	9
3.	Command Overhead	10
4.	Mechanical Positioning Performance	11
5.	Full Stroke Seek Time	12
6.	Cylinder Skew	12
7.	Head Skew	12
8.	Latency Time	12
9.	Drive Ready Time	13
10.	Data Transfer Speed	13
11.	Simple Sequential Access Performance	14
12.	Random Access Performance	14
13.	Mode Transition Time	15
14.	Power Connector Pin Assignments	21
15.	Table of signals	22
16.	Logic Signal Levels	23
17.	J1 Jumper Pins Location	23
18.	J1 Pin Assignment	24
19.	Default Jumper Setting	25
20.	Input Voltage	27
21.	Power Supply Current (with the SCSI termination power enabled)	27
22.	Power Supply Generated Ripple as seen at file power connector	27
23.	Outline of DAL5-3540	29
24.	Physical Dimension	29
25.	Mechanical Dimension	30
26.	Connector Locations	31
27.	Connector and Jumper Location	32
28.	Mounting Positions and the Tappings	33
29.	Random Vibration PSD Profile Breakpoints (Operating)	35
30.	Random Vibration PSD Profile Breakpoints (Non-Operating)	35
31.	A-weighted Sound Power Levels	37
32.	SCSI Commands Supported	43
33.	SCSI Commands Supported	44
34.	FORMAT UNIT (04h)	46
35.	Format of Defect List Header	47
36.	Format of Defect List	48
37.	INQUIRY (12h)	49
38.	INQUIRY DATA Invalid LUN Specified	51
39.	INQUIRY Data - EVPD = 0	52
40.	Product ID vs. Formatted Capacity	53
41.	INQUIRY DATA when EVPD = 1 (Page Code = 00h)	55
42.	INQUIRY DATA when EVPD = 1 (Page Code = 01h)	56
43.	INQUIRY DATA - EVPD = 1 (Page Code = 03h)	57
44.	INQUIRY DATA - EVPD = 1 (Page Code = 80h)	58
45.	INQUIRY DATA - EVPD = 1 (Page Code = 82h)	59
46.	MODE SENSE (1Ah)	62
47.	Page Code Usage	63
48.	MODE Parameter List (Header)	64
49.	MODE Parameter Block Descriptor	65
50.	MODE Parameter Page Descriptor	66
51.	MODE SELECT (15h)	67

52.	Page Code Usage	69
53.	Page 0h	70
54.	Page 1h	73
55.	Page 2h	78
56.	Page 3h	79
57.	Page 4h	81
58.	Page 7h	82
59.	Page 8h	84
60.	Page Ah	86
61.	Page 0Dh	88
62.	Pre-Fetch (34h)	89
63.	READ (08h)	91
64.	READ CAPACITY (25h)	92
65.	Format of READ CAPACITY command reply	93
66.	Read Defect Data (37h)	94
67.	Defect List Header	95
68.	Defect Descriptors of Bytes from Index Format	95
69.	Defect Descriptor of Physical Sector Format	96
70.	Read Extended (28h)	97
71.	READ BUFFER (3Ch)	98
72.	READ BUFFER Header	99
73.	READ BUFFER DESCRIPTOR	100
74.	READ LONG (3Eh)	101
75.	REASSIGN BLOCKS (07h)	102
76.	Format of REASSIGN BLOCKS data	103
77.	RECEIVE DIAGNOSTIC (1Ch)	104
78.	RELEASE (17h)	106
79.	REQUEST SENSE (03h)	107
80.	RESERVE (16h)	108
81.	REZERO UNIT (01h)	110
82.	SEEK (0Bh)	111
83.	SEEK EXTENDED (2Bh)	112
84.	SEND DIAGNOSTIC (1Dh)	113
85.	START/STOP Unit (1Bh)	116
86.	SYNCHRONIZE CACHE (35h)	117
87.	TEST UNIT READY (00h)	118
88.	VERIFY (2Fh)	119
89.	WRITE (0Ah)	120
90.	WRITE EXTENDED (2Ah)	121
91.	WRITE AND VERIFY (2Eh)	122
92.	WRITE BUFFER (3Bh)	123
93.	WRITE BUFFER Header	124
94.	WRITE LONG (3Fh)	127
95.	SCSI Status Byte	129
96.	Supported Messages	131
97.	Synchronous Data Transfer Request.	132
98.	Initiator Request/Target Response	133
99.	Target Response to Initiator's Transfer Period	133
100.	Queue Tag Messages	136
101.	Read Ahead handling per each command	163
102.	Format of Sense Data	173
103.	Field Pointer Bytes	180
104.	Actual Retry Count Bytes	181
105.	Format Progress Indication Bytes	181
106.	Read Recovery Steps.	186

107. Write Recovery Steps. 187
108. Servo Recovery Steps. 187
109. Spindle start Recovery Steps. 187

1.0 General

1.1 Introduction

This paper describes functional and interface specifications for 3.5" hard disk drives of DAL5-3540.

The interface conforms to the 'draft proposed American National Standard for information systems - SMALL COMPUTER SYSTEM INTERFACE - 2 (SCSI-2) October 17, 1991'.

The vendor specific items and options supported by the drive are described in each section.

1.2 References

- Draft proposed American National Standard for information systems - SMALL COMPUTER SYSTEM INTERFACE - 2 (SCSI-2) October 17, 1991
-

1.3 Glossary

<i>Word</i>	<i>Meaning</i>
Kbpi	1,000 Bits Per Inch
Mbps	1,000,000 Bits per second
MB	1,000,000 bytes
KB	1,000 bytes unless otherwise specified
32KB	32 x 1,024 bytes
64KB	64 x 1,024 bytes
Mb/sq.in	1,000,000 bits per square inch
MLC	Machine Level Control

1.4 General Caution

The drive can be easily damaged by shocks or ESD (Electric Static Discharge), so any damages applied to the drive after taking out from shipping package and opening ESD protective bag are user's responsibilities.

.

1.0 Outline of the drive

- SCSI 2 Standard
- Fast SCSI (up to 10 Mbytes/sec transfer)
- Queued I/O Process support
- 512 bytes/sector sector format
- Interleave factor 1:1
- Write cache for sequential write operation
- 4 x 16 KBytes segmented sector buffer
- Advanced ECC on the fly (EOF)
- Automatic error recovery procedures for read and write commands
- Self diagnostics on power on and resident diagnostics
- Transparent defect management with ADR(Automatic Defect Reallocation) during write cache.
- 12 msec seek time in read operation.
- 4500rpm spindle rotation.
- Closed loop actuator servo
- Dedicated head landing zone
- Automatic actuator lock

Part 1. Functional Specification

2.0 Fixed Disk Subsystem Description

2.1 Control Electronics

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

- Conducts a power-up sequence and calibrates the servo.
- Monitors various timers for head settling, servo failure, etc.
- Analyzes servo signals to provide closed loop control. These include position error signal and estimated velocity.
- Controls the voice coil motor driver to align the actuator onto a desired position.
- Monitors the actuator position and determines the target track for a seek operation.
- Constantly monitors error conditions of the servo and takes corresponding action if an error occurs.
- Controls starting, stopping, and rotating speed of the spindle.
- Controls and interprets all interface signals between the host controller and the drive.
- Controls read write accessing of the disk media, including defect management and error recovery.
- Performs self-checkout (diagnostics).

2.2 Head Disk Assembly

The head disk assembly (HDA) is assembled in a clean room environment and contain a disk and actuator assembly. Air is constantly circulated and filtered when the drive is operational. Venting of the HDA is accomplished via a breather filter.

The spindle is driven directly by a brushless, sensorless DC drive motor. Dynamic braking is used to stop the spindle quickly.

2.3 Actuator

The read/write heads are mounted in the actuator. The actuator is a swing-arm assembly driven by a voice coil motor. A closed-loop positioning servo controls the movement of the actuator. An embedded servo pattern supplies feedback to the positioning servo to keep the read/write heads centered over the desired track.

The actuator assembly is balanced to allow vertical or horizontal mounting without adjustment.

When the drive is powered off, the actuator automatically moves the head to a dedicated landing zone outside of the data area, where the actuator is locked.

3.0 Drive Characteristics

This chapter provides the characteristics of the drive.

3.1 Formatted Capacity

Description	DALS-3540	
Label Capacity (MB)	540	
Bytes per Sector	512	
Sectors per Track	84-128	
Number of heads	2	
Number of disks	1	
Spare Sectors per Zone	0	
Data Cylinder per Zone	382-1023	
Total Physical Data Bytes	550216704	
Number of LBAs	1056768	
Total Logical Data Bytes	541065216	

Figure 1. Formatted Capacity

3.2 Data Sheet

Figure 2. Data Sheet	
Buffer to/from media [Mbit/sec]	32.25 to 48.75
Host to/from buffer (interface transfer rate) [Mbyte/sec]	10
Data buffer size	64Kbyte
Number of buffer segments	4 x 16Kbyte
Rotational speed [RPM]	4500
Recording density [Kbpi]	68.4(Ave) / 82.6(Max)
Track density [TPI]	5340
Areal density [Mb/sq.in.]	365(Ave) / 441(Max)
Data bands	8

3.3 Performance Characteristics

A file performance is characterized by the following parameters:

- Command Overhead
- Mechanical Positioning
 - Seek Time
 - Latency
- Data Transfer Speed
- Buffering Operation (Lookahead/Write cache)

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

3.3.1 Command Overhead

Command overhead is defined as the time required:

- From the time that the drive is selected
- to the time available for the first data byte of a READ command when the requested data is not in the buffer
- exclude
 - Physical seek time
 - Latency time
 - SCSI BUS handling delay by host

Command Case (File is in quiescence state)	Time
Read (Cache not hit)	< 1.0 [msec]
Read (Cache hit)	< 0.8 [msec]
Write	< 0.8 [msec]
Seek	< 2.0 [msec]

Note: The above table gives an average time.

3.3.2 Mechanical Positioning

3.3.2.1 Average Seek Time (Including Settling)

Figure 4. Mechanical Positioning Performance		
Command Type	Typical	Max
Read	12 [msec]	13 [msec]
Write	12.7 [msec]	14 [msec]

“Typical” and “Max” are given throughout the performance specification by;

Typical Average of the drive population tested at nominal environmental and voltage conditions.
Max Maximum value measured on any one drive over the full range of the environmental and voltage conditions. (See 6.2, “Environment” on page 26 and 6.3, “DC Power Requirements” on page 27 for ranges.)

The seek time is measured from the start of actuator's motion to the start of **a reliable read or write operation**. Reliable read or write implies that error correction/recovery is not used to correct arrival problems. The average seek time is measured as the weighted average of all possible seek combinations.

$$\text{Weighted Average} = \frac{\sum_{n=1}^{\max} (\max + 1 - n) (T_{n.in} + T_{n.out})}{(\max + 1) (\max)}$$

Where:

max = Maximum seek length

n = Seek length (1 to max)

T_{n.in} = Inward measured seek time for an n track seek

T_{n.out} = Outward measured seek time for an n track seek

3.3.2.2 Full Stroke Seek

Figure 5. Full Stroke Seek Time		
Function	Typical	Max.
Read [msec]	20	25
Write [msec]	21	27

Full stroke seek is measured as the average of 1000 full stroke seeks with a **random head switch** from both directions (inward and outward).

3.3.2.3 Cylinder Switch Time (Cylinder Skew)

Figure 6. Cylinder Skew	
	Typical
Cylinder Skew	4.4 [msec]

A cylinder switch time is defined as the amount of time required by the fixed disk access the next sequential block after reading the last sector in the current cylinder.

The measured method is given in 3.3.6, “Throughput” on page 14.

3.3.2.4 Head Switch Time (Head Skew)

Figure 7. Head Skew	
	Typical
Head Skew	3.5 [msec]

3.3.2.5 Average Latency

Figure 8. Latency Time		
Rotation	Time for a revolution	Average Latency
4500 [RPM]	13.3 [msec]	6.67 [msec]

3.3.3 Drive Ready Time

Figure 9. Drive Ready Time		
Condition	Typical	Max.
Power On to Ready	8 [sec]	31 [sec]

Ready The condition in which the drive is able to perform a media access command (eg. read, write) immediately.

Power On This includes the time required for the internal self diagnostics.

3.3.4 Data Transfer Speed

Figure 10. Data Transfer Speed	
Description	Typical
Disk-Buffer Transfer (Zone 0)	
(Instantaneous)	4.92 [Mbyte/sec]
(Sustained)	3.74 [Mbyte/sec]
Disk-Buffer Transfer (Zone 7)	
(Instantaneous)	3.23 [Mbyte/sec]
(Sustained)	2.52 [Mbyte/sec]
Buffer-Host	
FAST SCSI	10 [Mbyte/sec]

- Instantaneous disk-buffer transfer rate (Mbyte/sec) is derived by:
 $(\text{Number of sectors on a track}) * 512 * (\text{revolution/sec})$

Note: Number of sectors per track will vary because of the linear density recording.

- Sustained disk-buffer transfer rate (Mbyte/sec) is defined by considering head/cylinder change time. This gives a local average data transfer rate. It is derived by:

$$(\text{Sustained Transfer Rate}) = A / (B + C + D)$$

$$A = (\text{Number of data sectors per cylinder}) * 512$$

$$B = ((\text{\# of Surface per cylinder}) - 1) * (\text{Head switch time})$$

$$C = (\text{Cylinder change time})$$

$$D = (\text{\# of Surface}) * (\text{One revolution time})$$

- Instantaneous Buffer-Host Transfer Rate (Mbyte/sec) defines the maximum data transfer rate on AT Bus. It also depends on the speed of the host.

The measurement method is given in 3.3.6, "Throughput" on page 14.

3.3.5 Buffering Operation (Lookahead/Write Cache)

In order to improve the total performance, the file utilizes it's own buffer for lookahead. The total of 64K bytes of the buffer is divided into two segmented blocks. These two segments are managed intelligently to obtain the best performance.

3.3.6 Throughput

3.3.6.1 Simple Sequential Access

Figure 11. Simple Sequential Access Performance		
Operation	Typical	Max
Sequential Read/Write	(4.69 [sec])	(4.91 [sec])

The above table gives the time required to read/write for a total of 8000x consecutive blocks (16,777,216 bytes) accessed by 128 read/write commands. Typical and Max values are given by 105% and 110% of T respectively throughout following performance description.

Note: Assumes a host system responds instantaneously.

$$T = (A * 128) + B + C + 16,777,216/D + 512/E \quad (\text{READ})$$

$$T = (A * 128) + B + C + 16,777,216/D \quad (\text{WRITE})$$

where:

T = Calculated Time (sec)

A = Command Process Time (Pre/Post Command overhead)

B = Average Seek Time

C = Average Latency

D = Sustained Disk-Buffer Transfer Rate (Mbyte/sec)

E = Buffer-Host Transfer Rate (Mbyte/sec)

3.3.6.2 Random Access

Figure 12. Random Access Performance		
Operation	Typical	Max
Random Read	82 [sec]	88 [sec]
Random Write	82 [sec]	90 [sec]

The above table gives the time required to execute a total of 1000x read/write commands which access a random LBA.

$$T = (A + B + C + 512/D + 512/E) * 4096 \quad (\text{READ})$$

$$T = (A + B + C + 512/D) * 4096 \quad (\text{WRITE})$$

where:

T = Calculated Time (sec)

A = Command Process Time (Pre/Post Command overhead)

B = Average Seek Time

C = Average Latency

D = Sustained Disk-Buffer Transfer Rate (Mbyte/sec)

E = Buffer-Host Transfer Rate (Mbyte/sec)

3.3.7 Operating Mode Definition

Operating Mode	Description
Spin-Up	Start up time period from spindle stop or power down.
Seek	Seek operation mode
Write	Write operation mode
Read	Read operation mode
Idle	Spindle motor and servo system are working normally. Commands can be received and processed immediately.
Standby	Spindle motor is stopped. Commands can be received immediately, but write or read operations cannot begin until the spindle is spun-up and the Servo system is ready.

Note: Upon Power down or Spindle stopped, a head locking mechanism will secure the heads in the ID parking position.

Recovering from Standby mode does not need soft reset nor hard reset.

3.3.7.1 Mode Transition Time

Figure 13. Mode Transition Time			
From	To	Typical	Max
Standby	Idle	8 [sec]	31 [sec]
Idle	Standby	Immediate	N/A

Note: The actual spin down time will exist, however the command will be processed immediately.

4.0 Data integrity

The drive retains recorded information under all non-write operations.

No more than one sector can be lost by power down during write operation while write cache is disabled. In case that power down occurs before completion of data transfer from write cache to disk while write cache is enabled, the data remaining in write cache can be lost. To prevent this data loss at power off, one of following actions are recommended:

- SYNCHRONIZE CACHE (35h) command is completed.
- Write cache of Page-8h is disable by MODE SELECT command before the write access.

4.1 Equipment Status

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

- Access recalibration/tuning is complete.
- Spindle speed meets requirements for reliable operations.
- Self-check of drive is complete.

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

- Spindle speed goes outside of requirements for reliable operation.
- “Write fault” is detected.

4.2 Error Recovery

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

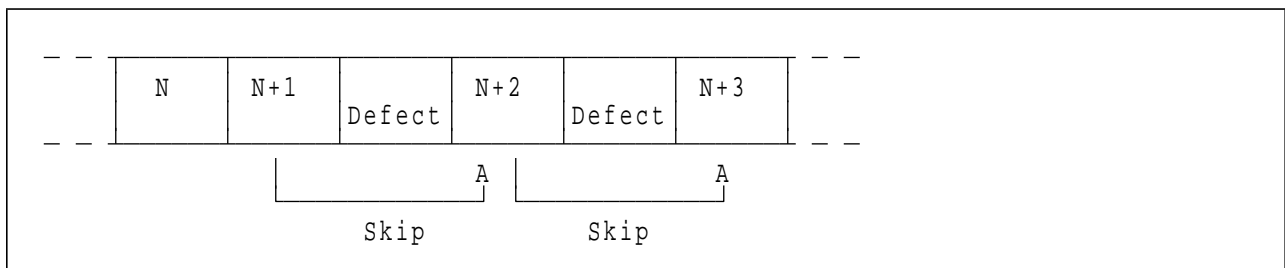
Errors that are uncorrectable after application of the error recovery procedures are reported to the host system as non-recoverable errors.

5.0 Physical Format

Media defects are remapped to the next available sector during Format Process in manufacturing. The mapping from LBA to the physical locations is calculated by an internal maintained table.

5.1 Shipped Format

- Data areas are optimally used.
- No extra sector is wasted as a spare throughout user data areas.
- All pushes generated by defects are absorbed by spare tracks of inner zone.



Defects are skipped without any constraint, such as track or cylinder boundary. The calculation from LBA to physical is done automatically by internal table.

6.0 Specification

6.1 Electrical interface specification

6.1.1 Connectors

6.1.1.1 Power

The DC power connector is designed to mate with AMP (part 1-480424-0) using AMP pins (part 350078-4) strip or (part 61173-4) loose piece, or their equivalents. Pin assignments are shown below.




Figure 14. Power Connector Pin Assignments

6.1.1.2 SCSI Signal Connector

The SCSI signal connector is 50-pin connector meeting the ANSI SCSI specification.

6.1.2 Signal Definition

The pin assignments of interface signals are listed as follows:

Figure 15. Table of signals

PIN	SIGNAL	PIN	SIGNAL
01	Ground	02	-DB(0)
03	Ground	04	-DB(1)
05	Ground	06	-DB(2)
07	Ground	08	-DB(3)
09	Ground	10	-DB(4)
11	Ground	12	-DB(5)
13	Ground	14	-DB(6)
15	Ground	16	-DB(7)
17	Ground	18	-DB(P)
19	Ground	20	Ground
21	Ground	22	Ground
23	Ground	24	Ground
25	Open	26	TRM Power
27	Ground	28	Ground
29	Ground	30	Ground
31	Ground	32	-ATN
33	Ground	34	Ground
35	Ground	36	-BSY
37	Ground	38	-ACK
39	Ground	40	-RST
41	Ground	42	-MSG
43	Ground	44	-SEL
45	Ground	46	-C/D
47	Ground	48	-REQ
49	Ground	50	-I/O

6.1.3 SCSI Cable

The drive uses single-ended drivers and receivers which will permit cable lengths of up to 6 meters (19.68 feet).

For a single ended cable a 50 conductor flat cable or a 25 signal twisted cable can be used with a maximum length of 6.0 meters, and a stub length not exceeding 0.1 meter.

6.1.4 SCSI Bus Terminator

The drive has an internal Active SCSI bus terminator, and can be controlled on/off with one jumper block provided at the card edge. The users of the drive are responsible for proper termination and power source for SCSI bus.

6.1.5 SCSI Bus Electrical Characteristics

The interface logic signals has the following electrical specifications:

Inputs	:	Input High Voltage	–	2.0 V min.
		Input Low Voltage	–	0.8 V max.
Outputs	:	Output High Voltage	–	High Impedance
		Output Low Voltage	–	0.5 V max./IO=48mA current

Figure 16. Logic Signal Levels

The 12 positions jumper block shown below is to select the SCSI device ID, Auto spin up option control, Unit Attention option control, SCSI terminator ON/OFF setting, and T.I. Sync Negotiation option control.




Figure 17. J1 Jumper Pins Location

6.1.6 The pin assignment

Figure 18. J1 Pin Assignment

Pin#	Name	Description	Function (if jumper connected)	Default (Note)
1	-DAS0	SCSI ID Selct #0	SCSI ID Bit 0	OFF
2	GND	GND		
3	-DAS1	SCSI ID Selct #1	SCSI ID Bit 1	ON
4	GND	GND		
5	-DAS2	SCSI ID Selct #2	SCSI ID Bit 2	ON
6	GND	GND		
7	-ASD	Auto Spin-up disable	Need START COMMAND to start spindle during Power On Reset (POR).	OFF
8	GND	GND		
9	-UAD	Unit Attention disable	Unit Attention Condition will NOT be posted.	OFF
10	GND	GND		
11	-TERM_ON	SCSI Terminator ON	Connect to Term Power. (Terminator ON)	ON
12	GND	GND		
13	-TISND	Sync Negotiation Disable	Target will NOT initiate Synchronous Negotiation.	ON
14	GND	GND		
15	DLY_SPN	Spare	Reserve	OFF
16	GND	GND		
17	HAN5	Spare	Reserve	OFF
18	GND	GND		
19	HAN6	Spare	Reserve	OFF
20	GND	GND		
21	SYNC	Spare	Reserve	OFF
22	GND	GND		
23	-MYBUSY	SCSI Busy (Connected)	Use to drive an external LED. Up to 8mA of sink current capability is provided.	OFF
24	5 V	5 Volt (Connected to +5V through 510 ohm resistor)		

Note: "Default" indicates the shipping default jumper setting.
ON indicates that the jumper is connected.
OFF indicates that the jumper is not installed.

6.1.7 Signal description

Throughout this paragraph ON means a shunt jumper is installed, and OFF means that no shunt jumper is installed.

6.1.7.1 Device address select lines.(-DAS0, -DAS1, -DAS2)

These three lines defines DAL5-3540 device ID on the SCSI BUS. -DAS0 is the least significant bit, and -DAS2 is the most significant bit. Device ID is defined as follows.

1-2	3-4	5-6	Device ID	
off	off	off	0	
on	off	off	1	
off	on	off	2	
on	on	off	3	
off	off	on	4	
on	off	on	5	
off	on	on	6	<—————Shipping default ID
on	on	on	7	

Figure 19. Default Jumper Setting

6.1.7.2 GND pins.(Pin #2,4,6,8,10,12,14,16,18,20,22)

Connected to signal ground.

6.1.8 Shipping Default Condition

The drive will be shipped with the following default condition.

- SCSI ID = 6 (shunt jumpers are installed at Pin#3-4 & #5-6)
- Auto Spindle enabled (No shunt jumper is installed at Pin#7-8)
- Unit Attention enabled (No shunt jumper is installed at Pin#9-10)
- SCSI Terminator ON (A shunt jumper is installed at Pin#11-12)
- TI Sync Negotiation disabled (A shunt jumper is installed at Pin#13-14)

6.1.8.1 Mechanical Outline

The card with disk enclosure mechanical outline is shown in Figure 27 on page 32.

6.2 Environment

6.2.1 Temperature and Humidity

Operating Conditions	
Temperature	5 to 55[°C] (See note)
Relative Humidity	8 to 90 [% RH] non-condensing
Maximum Wet Bulb Temperature	29.4[°C] non-condensing
Maximum Temperature Gradient	15[°C / Hour]
Altitude	- 300 to 3048 [m]
Non-Operating Conditions	
Temperature	- 40 to 65[°C]
Relative Humidity	5 to 95 [% RH] non-condensing
Maximum Wet Bulb Temperature	35[°C] non-condensing
Maximum Temperature Gradient	15[°C / Hour]
Altitude	- 300 to 12,000 [m]
Note: The system has to provide sufficient ventilation to maintain a surface temperature below 60[°C] at the center of the top cover of the drive.	

6.3 DC Power Requirements

The following voltage specification apply at the file power connector. Damage to the file electronics may result if the power supply cable is connected or disconnected while power is being applied to the file (No hot plug/unplug is allowed). There is no special power on/off sequencing required.

Figure 20. Input Voltage		
	During run and spin up	Absolute max voltage
+ 5 Volts Supply	5V +/- 5%	7V
+12 Volts Supply	12V +10% , - 8%	15V

Figure 21. Power Supply Current (with the SCSI termination power enabled)					
(All values in Amps.)	+5Volts		+12Volts		Total (W)
	Pop Mean	Std.Dev	Pop Mean	Std.Dev	
Idle Average	0.23	0.02	0.12	0.02	2.6
Idle ripple (peak-to-peak)	0.21	0.02	0.11	0.02	
Seek peak (*1)	0.61	0.03	0.52	0.03	
Seek average (*1)	0.41	0.02	0.27	0.02	5.2
Start up (max)	0.51	0.02	1.16	0.03	
RND R/W peak (*2)	0.75	0.03	0.50	0.03	
RND R/W average (*2)	0.57	0.03	0.14	0.02	4.5
Standby/Sleep average	0.20	0.01	0.01	0.01	1.1

Notes:

1. Random Seeks at 100% duty cycle.
2. Seek Duty = 30%, W/R Duty = 45%, Idle Duty = 25%.

Figure 22. Power Supply Generated Ripple as seen at file power connector		
	Maximum	Notes
+5 V DC	100 [mV pp]	0-10 [MHz]
+12 V DC	150 [mV pp]	0-10 [MHz]

During file start up and seeking, 12 volt ripple is generated by the file (referred to as dynamic loading). If several files have their power daisy chained together then the power supply ripple plus other file's dynamic loading must remain within the regulation tolerance of +10/-8%. A common supply with separate power leads to each file is a more desirable method of power distribution.

To prevent external electrical noise from interfering with the file's performance, the file must be held by four screws in a user system frame which has no electrical level difference at the four screws position, and has less than +/-300 milivolts peak to peak level difference to the file power connector ground.

6.4 Reliability

6.4.1 Mean Time Between Failures (MTBF)

350,000 power-on hours (POH) under the following conditions:

- 4,000 POH per year (333 POH per month)
- Seeking/writing/reading operation to be 20% of POH at 40°C or lower environmental temperature.

6.4.2 Contact Start Stop (CSS)

The drive withstands a minimum of 40,000 contact start/stop cycles under 40°C.

6.4.3 Useful Life

The useful life of the drive is 5 years minimum.

6.4.4 Preventive Maintenance

None.

6.4.5 Data Reliability

- Probability of not recovering data 1 in 10¹³ bits read
- ECC implementation

A Read Solomon Error Code of degree-8 with non-interleaved is used to cover the data field. The ECC polynomial is derived from

$$g(X) = (X + 1)(X + A)(X + A^2) \dots (X + A^7)$$

On-The-Fly correction implemented in the file covers four or less symbols of error in one sector. (One symbol is 10 bits.)

6.4.6 Seek/ID Mis-compare Errors

A non-recoverable seek/ID mis-compare error is defined as a seek operation that cannot be corrected by fixed disk error recovery procedure. Seek errors occurring for field format operations are considered to be non-recoverable.

No drive has more than one non-recoverable seek/ID mis-compare error per 5 million seek operations (1 in 5 x 10⁶) when operated at the full range of voltage and environmental conditions.

Non-recoverable seek/ID mis-compare errors indicate a defective drive.

6.4.7 Equipment Errors

A recoverable equipment error is any error other than a seek/ID mis-compare error or read error that is detected and corrected by the drive error recovery procedure. Examples are Write Fault, Drive Not Ready and internal drive errors.

No drive has more than one recoverable equipment error per 10⁸ reads, 10⁶ writes or 10⁶ seeks operations when operated at the full range of voltage and environmental conditions.

Non-recoverable equipment errors indicate a defective drive.

6.5 Mechanical Specifications

6.5.1 Outline




Figure 23. Outline of DALS-3540

6.5.2 Mechanical Dimensions

The following chart describes the dimensions and the weight.

Figure 24. Physical Dimension	
Height [mm]	25.4 ± 0.4
Width [mm]	101.6 ± 0.4
Length [mm]	146.0 ± 0.6
Weight [gram]	450 Max.




Figure 25. Mechanical Dimension

6.5.3 Connector Locations




Figure 26. Connector Locations

6.5.4 Connector and Jumper

Jumpers are used for the settings of SCSI IDs, Auto Spin Up, Unit Attention, Active Terminator and Target Initiated Synchronous Negotiation.

6.5.4.1 Connector and Jumper Location




Figure 27. Connector and Jumper Location

6.5.4.2 Mounting Positions and the Tappings




Figure 28. Mounting Positions and the Tappings

6.5.5 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 0.6 - 1.0 [Nm] (6 - 10 [Kgf.cm]). The recommended mounting screw depth is 6 [mm] Max for bottom and 3.5 [mm] Max for horizontal mounting.

For reliable operation, it is recommended to mount the drive in the system securely enough to prevent from excessive motion or vibration of the drive at seek operation or spindle rotation, with using appropriate screws or equivalent mounting hardware. Consult with the issuer of this specification for actual application if necessary.

Drive level vibration test and shock test are to be conducted with mounting the drive to the table using bottom four screws.

6.5.6 Shipping Zone and Lock

A dedicated "shipping" (or "landing") zone on the disk, not on the data area of the disk, is provided to protect the disk data during shipping, movement, or storage. Upon power down, a heads are automatically parked and a head locking mechanism will secure the heads in this zone.

6.6 Vibration and Shock

All vibration and shock measurements in this section are made with the drive that has no mounting attachments for the systems. The input power for the measurements is applied to the normal drive mounting points.

6.6.1 Operating Vibration

The drive operates with no unrecoverable errors while being subjected to the following vibration levels.

The measurements are carried out during 30 minutes of random vibration using the power spectral density (PSD) levels specified in IBM standards as V5L. The vibration test level for V5L is 0.67G (RMS).

Figure 29. Random Vibration PSD Profile Breakpoints (Operating)									
Hz	Random Vibration PSD Profile Breakpoints (Operating)								
[Hz]	5	17	45	48	62	65	150	200	500
$\times 10^{-3}$ [G ² /Hz]	0.02	1.1	1.1	8.0	8.0	1.0	1.0	0.5	0.5

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

6.6.2 Non-Operating Vibrations

The drive does not sustain permanent damage or loss of recorded data after being subjected to the environment described below.

6.6.2.1 Random Vibration

The test consists of a random vibration applied for each of three mutually perpendicular axes with the time duration of 15 minutes per axis. The PSD levels for the test simulates the shipping and relocation environment which is shown below.

Figure 30. Random Vibration PSD Profile Breakpoints (Non-Operating)							
Hz	Random Vibration PSD Profile Breakpoints (Non-Operating)						
Hz	2	4	8	40	55	70	200
[G ² /Hz]	0.001	0.03	0.03	0.003	0.01	0.01	0.001

Overall RMS (Root Mean Square) level of vibration is 1.04G (RMS).

6.6.2.2 Swept Sine Vibration

- 2 G (Zero to peak), 5 to 500 to 5 Hz sine wave
- 0.5 oct/min sweep rate
- 3 minutes dwell at two major resonances

6.6.3 Operating Shock

The drive meets the following criteria.

- No errors within shock pulses of 5G, 11ms half-sine wave.
- No data loss, seek errors, or permanent damages within shock pulses of 10G, 11 ms half-sine wave.
- No data loss or permanent damages at Idle, Seek and Read modes within shock pulses of 30G 4ms, or 15G 5ms half-sine wave.

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

6.6.4 Non-Operating Shock

The drive withstands without damage or degradation of performance, a 75G half-sine wave shock pulse of 11 ms duration on six sides when heads are secured. (When the power is not applied to the unit, the heads are automatically located on a landing zone.)

All shocks are applied for each direction of the drive for three mutually perpendicular axes and one axis at a time. Input levels are measured on a base plate where the drive is attached with four screws.

6.7 Acoustics

The following shows the acoustic levels.

6.7.1 Sound Power Levels

The upper limit criteria of the A-weighted sound power levels are given in Bel relative to one pico watt and are shown in the following table.

Figure 31. A-weighted Sound Power Levels	
Mode	A-weighted Sound Power Level [Bel]
Idle	4.5
Operating	4.8

Background power levels of the acoustic test chamber for each octave band are to be recorded.

Sound power levels are measured with the drive supported by spacers so that the lower surface of the drive is located 25 ± 3 mm height from the chamber desk. No sound absorbing material shall be used.

The acoustical characteristics of the drive subsystem are measured under the following conditions.

Idle mode:

Powered on, disks spinning, track following, unit ready to receive and respond to control line commands.

Operating mode:

Continuous random cylinder selection and seek operation of actuator with a delay for a time period achieving the required seek rate N_s according to the following formula:

$$N_s = 0.4 / (T_t + T_l)$$

where:

N_s = average seek rate in seeks/sec.

T_t = published random seek time.

T_l = time for the drive to rotate by half a revolution.

6.7.2 Sound Power Acceptance Criteria

Statistical upper limit $(L_W)_{stat}$ is calculated with the following formula.

$$(L_W)_{stat} = (L_W)_m + k \times (s_t)_W$$

where:

$(L_W)_m$ is the mean value of the A-weighted sound power level for samples of N drives.

$(s_t)_W$ is the total standard deviation for A-weighted sound power level.

$$(s_t)_W = \text{SQRT}((s_R)_W^2 + (s_P)_W^2)$$

$(s_R)_W$ is the standard deviation for sound power level.

Assume $(s_R)_W = 0.075$ B.

$(s_P)_W$ is the standard deviation of the samples for A-weighted sound power level.

k is a coefficient determined by number of samples (N) as shown below.

N	3	4	5	6	7	8	9	10	11	12	13	14	15
k	3.19	2.74	2.74	2.49	2.33	2.22	2.13	2.07	2.01	1.97	1.93	1.90	1.87

The calculated left hand side of the criterion equation above is rounded to the nearest 0.05 bel. The individual terms may be rounded to the nearest 0.01 bel before calculation.

6.8 Identification Labels

The following labels are attached to the drive.

1. A label placed on the top of the HDA contains the statement “Made by IBM” or equivalent, Part number, and MLC number.
2. A bar code label placed on the drive is based on user requests. The location is to be designated in the drawing.
3. Labels containing the vendor's name, disk drive model number, serial number, place of manufacture and UL/CSA/TÜV certificates.

The labels may be integrated.

6.9 Electromagnetic Compatibility

The drive, when installed in the host system and exercised with a random accessing routine at maximum data rate, meets the worldwide EMC requirements listed below.

IBM will provide technical support to meet the requirements to comply with the EMC specifications.

- United States Federal Communications Commission (FCC) Rules and Regulations (Class B), Part 15. IBM Corporate Standard C-S 2-0001-026 (6 dB buffer should be maintained on the emission requirements).
- European Economic Community (EEC) directive number 76/889 related to the control of radio frequency interference and the Verband Deutscher Elektrotechniker (VDE) requirements of Germany (GOP).

6.10 Safety

The following shows the safety standards for different countries.

6.10.1 Underwriters Lab (UL) Approval

DALS-3540 comply with UL 1950.

6.10.2 Canadian Standards Authority (CSA) Approval

DALS-3540 comply with CSA C22.2 #950-M89.

6.10.3 IEC Compliance

DALS-3540 comply with IEC 380, IEC 435 and IEC 950.

6.10.4 German Safety Mark

DALS-3540 is approved by TÜV on Test Requirements: EN 60 950:1988/A1:1990/A2:1991.

6.10.5 Flammability

The printed circuit boards used in this product is made of material with the UL recognized flammability rating of V-1 or better. The flammability rating is marked or etched on the board. All other parts not considered electrical components are made of material with the UL recognized flammability rating of V-1 or better, except small mechanical parts.

6.10.6 Safe Handling

The products are conditioned for safe handling in regards to sharp edges and corners.

6.10.7 Environment

The product does not contain any known or suspected carcinogens.

Environmental controls meet or exceed all applicable government regulations in the country of origin. Safe chemical usage and manufacturing control are used to protect the environment. An environmental impact assessment has been done on the manufacturing process used to build the drive, the drive itself, and the disposal of the drive at the end of its life.

Production also meets the requirements of the international treaty on chlorofluorocarbon (CFC) control known as the United Nations Environment Program Montreal Protocol, and as ratified by the member nations. Materials to be controlled include CFC-11, CFC-12, CFC-113, CFC-114, CFC-115, Halon 1211, Halon 1301 and Halon 2402. Although not specified by the Protocol, CFC-112 is also controlled. In addition to the above protocol, IBM controls the following:

- All packaging materials used for the shipment of the product do not use controlled CFCs in the manufacturing process.
- All manufacturing processes for parts or assemblies including printed circuit boards, will not use the controlled CFC materials.

6.10.8 Secondary Circuit Protection

The drive uses printed circuit wiring that protects the possibility of sustained combustion due to circuit or component failure. Adequate secondary over-current protection is the responsibility of system suppliers.

The host system must protect the drive from any electrical short circuit problem. 10 [A] limit is required for safety purposes.

6.11 Packaging

The drives are packed in ESD protective bags for shipping.

Part 2. SCSI Interface Specification

7.0 SCSI COMMAND SET

Summaries of the SCSI commands supported by the drive are listed below. where O=Optional, M=Mandatory, E=Extended, R=Reserved and V=Vendor unique. The column of "SCSI-1" refers to ANSI version-1 standard. The column of "CCS" refers to the ANSI sub-committee Common Command Subset for DASD devices. The column "SCSI-2" refers to ANSI version-2 standard.

SCSI-1	CCS	SCSI-2	CODE	COMMAND
M	M	M	04h	FORMAT UNIT
E	M	M	12h	INQUIRY
O	O	O	4Ch	LOG SELECT
O	O	O	4Dh	LOG SENSE
O	O	O	15h	MODE SELECT
O	O	O	1Ah	MODE SENSE
O	O	O	34h	PRE-FETCH
M	M	M	08h	READ
R	O	O	3Ch	READ BUFFER
E	M	M	25h	READ CAPACITY
R	R	O	37h	READ DEFECT DATA
E	M	M	28h	READ EXTENDED
R	R	O	3Eh	READ LONG
O	O	O	07h	REASSIGN BLOCKS
O	O	O	1Ch	RECEIVE DIAGNOSTIC
O	M	M	17h	RELEASE
M	M	M	03h	REQUEST SENSE
O	M	M	16h	RESERVE
O	O	O	01h	REZERO UNIT
O	O	O	0Bh	SEEK
O	O	O	2Bh	SEEK EXTENDED
O	M	M	1Dh	SEND DIAGNOSTIC
O	O	O	1Bh	START/STOP UNIT
R	R	O	35h	SYNCHRONIZE CACHE
O	M	M	00h	TEST UNIT READY
O	O	O	2Fh	VERIFY
M	M	M	0Ah	WRITE
O	O	O	2Eh	WRITE AND VERIFY
R	O	O	3Bh	WRITE BUFFER
E	M	M	2Ah	WRITE EXTENDED
R	R	O	3Fh	WRITE LONG

Figure 32. SCSI Commands Supported. (In Alphabetical order)

SCSI-1	CCS	SCSI-2	CODE	COMMAND
O	M	M	00h	TEST UNIT READY
O	O	O	01h	REZERO UNIT
M	M	M	03h	REQUEST SENSE
M	M	M	04h	FORMAT UNIT
O	O	O	07h	REASSIGN BLOCKS
M	M	M	08h	READ
M	M	M	0Ah	WRITE
O	O	O	0Bh	SEEK
E	M	M	12h	INQUIRY
O	O	O	15h	MODE SELECT
O	M	M	16h	RESERVE
O	M	M	17h	RELEASE
O	O	O	1Ah	MODE SENSE
O	O	O	1Bh	START/STOP UNIT
O	O	O	1Ch	RECEIVE DIAGNOSTIC
O	M	M	1Dh	SEND DIAGNOSTIC
E	M	M	25h	READ CAPACITY
E	M	M	28h	READ EXTENDED
E	M	M	2Ah	WRITE EXTENDED
O	O	O	2Bh	SEEK EXTENDED
O	O	O	2Eh	WRITE AND VERIFY
O	O	O	2Fh	VERIFY
O	O	O	34h	PRE-FETCH
R	R	O	35h	SYNCHRONIZE CACHE
R	R	O	37h	READ DEFECT DATA
R	O	O	3Bh	WRITE BUFFER
R	O	O	3Ch	READ BUFFER
R	R	O	3Eh	READ LONG
R	R	O	3Fh	WRITE LONG
O	O	O	4Ch	LOG SELECT
O	O	O	4Dh	LOG SENSE

Figure 33. SCSI Commands Supported. (In Command Code order)

7.1 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 bit.

LSB. Least Significant bit.

7.2 Flag and Link Bits

Many of the structures in this section have fields names **FLAG** and **LINK**. The meaning of these fields is defined below.

FLAG The Flag bit specifies which message the target shall return to the initiator if the link bit is one and the command completes without any error. If Link is zero, Flag must be zero. If Link is one and the command terminates successfully, the drive will send either the **LINKED COMMAND COMPLETE** message (**FLAG=0**) or the **LINKED COMMAND COMPLETE WITH FLAG** message (**FLAG=1**). Typically this bit is used to cause an interrupt for 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. Upon successful completion of the command, the drive returns **INTERMEDIATE GOOD** status and then send one of the two messages defined under the Flag section shown above.

Upon unsuccessful completion of the command, the drive returns **Check Condition** status or **Reservation Conflict** status and then send the **COMMAND COMPLETE** message. No further commands in the chain are executed.

7.3 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 of four-byte field (0-byte, 1-byte, 2-byte, 3-byte), 0-byte is the most significant byte and 3-byte is the least significant byte.

7.4 FORMAT UNIT (04h)

	BIT							
	7	6	5	4	3	2	1	0
BYTE 0	Command code (04h)							
BYTE 1	LUN			Fmt Data	Cmp List	List Format		
BYTE 2	VU = 0							
BYTE 3	(MSB)							
BYTE 4	Interleave Factor (LSB)							
BYTE 5	VU = 0		RSVD = 0			Flag	Link	

Figure 34. FORMAT UNIT (04h)

The FORMAT command performs a physical formatting of the disk media. This includes handling of defective sectors, and the overwriting of all data areas with a constant data pattern. (Reserved areas of the media are not affected by the FORMAT command.)

- **FmtData** set to one specifies that a Data Out phase follows the Command phase. The Data Out phase consists of a defect list header followed by **zero defect descriptors**. FmtData set to zero specifies that no Data Out phase follows.
- **CmpList** set to one specifies that the GList (Grown Defect List) existing prior to the format is **not** used and is discarded. The Drive is formatted with PList and DList (if specified). DList becomes the new GList.

Note: The drive manages two internal defect lists and one external. The primary defect list (“P”List) is created at time of manufacture and cannot be altered. The grown defect list (“G”List) is built after manufacture by the Initiators by using of the REASSIGN BLOCK command and the Automatic Real-locate functions. The data defect list (“D”List) is an external list. It is supplied by an initiator in the DATA OUT phase of the FORMAT UNIT command.

- **List Format** specifies the format of the defect descriptor transferred to the Target when FmtData bit is set to one. This value must be set to zero as the only supported format is logical block. If the FmtData bit is set to zero this field must also be zero otherwise the command will complete with Check Condition with a Sense Key of illegal request and an additional sense code of invalid field in CDB.
- **Interleave Factor** may be zero or one, either of which specifies an interleave of 1:1. Other Interleave Factors are not supported because of the extensive buffering implemented in the drive.

7.4.1 Defect List

	BIT							
	7	6	5	4	3	2	1	0
BYTE 0	Reserved = 0							
BYTE 1	FOV	DPRY =0	DCRT	STPF	IP =0	DSP =0	Immd	0
BYTE 2	Defect list length MSB							
BYTE 3	Defect list length LSB							

Figure 35. Format of Defect List Header. Format of the defect list header sent during the data out phase when FmtData is set to one.

The Target has a limited implementation of the Format Option bits located in Bits 2 through 7 of Byte 1 of the Defect List Header (See Figure 35). If the Initiator attempts to select any function not implemented by the Target, the Target terminate 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 Parameter List.

- **FOV** (Format Options Valid) bit of zero causes the drive to verify that the setting for the DPRY (Disable Primary), DCRT (Disable Certification), STPF (stop Format), IP (Initialize Pattern), and DSP (Disable Saving Parameters) bits are zero. If any of these bits are not zero, the drive terminates 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 Parameter List.

Note: When FOV bit is one there is only one permissible combination of the DPRY, DCRT, STPF, IP and DSP bits. Any other combinations return Check Condition Status with a Sense Key of Illegal Request and an additional sense code of Invalid Field In Parameter List. The supported combination is:

DPRY=0 DCRT=1 STPF=1 IP=0 DSP=0

- **DPRY**(Disable Primary) bit set to zero indicates that the drive does not use portions of the medium identified as defective in the primary defect PList for Initiator addressable logical blocks. If the drive cannot locate the PList or it cannot determine whether a PList exists, the drive terminates the Format Unit command as described for STPF=1.
- **DCRT** (Disable certification) bit must be set to one. The drive does not generate a CList (certification list) nor perform a certification process while executing the Format Unit Command.
- **STPF** (stop Format) bit must be set to one. If one or both of the following conditions occurs, the drive terminates the Format Unit command with Check Condition Status. The sense key is set to Medium Error. and the additional sense code is set to either *Defect List Not Found* if the first condition occurred, or *Defect List Error* if the second condition occurred.
 - The Target cannot locate a required defect list nor determine that the list exists.
 - The Target encounters an unrecoverable error while accessing a required defect list.
- **IP** (Initialization Pattern) bit must be set to zero. The drive initializes all data with zeros.
- **DSP** (Disable Saving Parameters) bit must be set to zero. The drive saves all the Mode Select savable parameters during the format operation.
- **IMMD** (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. *Good Status* is returned following the CDB validation and transfer of data in the Data Out phase. If the immediate format operation terminates in error, Deferred Error Sense data is generated. With the immediate bit set to one, the Link bit must be set to zero.

The Defect List Length field specifies the total length in bytes of the defect descriptors that follow. The drive has an implementation limitation for number of defect descriptors. The number of defect descriptor shall be less than **128**. The defect list length must be equal to four times the number of defect descriptors to follow, otherwise the command is terminated with *Check Condition Status*. The sense key is set to *Illegal Request* and the additional sense code is set to *Invalid Field In Parameter List*. The defect descriptors must specify the defect based on the current Format Device parameters reported by the Mode Sense command.

The only format for the defect list supported by the drive is by logical block where the location of defective sectors is given by their LBA.

	BIT							
	7	6	5	4	3	2	1	0
BYTE 0	Defect Descriptor 0 LBA MSB							
BYTE 1	Defect Descriptor 0 LBA							
BYTE 2	Defect Descriptor 0 LBA							
BYTE 3	Defect Descriptor 0 LBA LSB							
BYTE 4n – 4n + 3	Defect Descriptor n LBA							

Figure 36. Format of Defect List. Format of the defect list sent during the data out phase when FmtData set to one.

7.5 INQUIRY (12h)

	BIT							
	7	6	5	4	3	2	1	0
BYTE 0	Command Code = 12h							
BYTE 1	LUN			RSVD = 0			EVPD	
BYTE 2	PAGE CODE							
BYTE 3	RSVD = 0							
BYTE 4	ALLOCATION LENGTH							
BYTE 5	VU = 0		RSVD = 0			FLAG	LINK	

Figure 37. INQUIRY (12h)

The INQUIRY command requests the parameters of the target (drive) to be sent to initiators.

An EVPD bit of one specifies that the drive shall return the optional vital product data specified by the code page field. **Page code** specifies which page of vital product data information the drive shall return.

EVPD	PAGE CODE	Description
0	0	The drive returns the standard INQUIRY data.
0	Non Zero	The drive returns CHECK CONDITION status with Sense Key of ILLEGAL REQUEST and Additional Sense Code of INVALID FIELD IN CDB.
1	Supported	The drive returns the vital product data of specified page code.
1	Non-supported	The drive returns CHECK CONDITION status with the sense key of ILLEGAL REQUEST and the additional sense code of INVALID FIELD IN CDB.

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

If an INQUIRY command is received from an initiator with a pending unit attention condition (before the the drive reports Check Condition status), the drive processes the INQUIRY command. The unit attention condition is not cleared by this action.

Two different formats for the INQUIRY data are defined.

- The first format is returned when an invalid LUN is specified by initiators.
- The second format is returned when a valid LUN is specified by initiators.

Each of these formats is described in the following sections.

Note: 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 alpha-numeric and represented in ASCII, not EBCDIC.

7.5.1 INQUIRY Data Format (When Invalid LUN is Specified)

BYTE	BIT							
	7	6	5	4	3	2	1	0
0	Qualifier			Peripheral Dev. Type=1Fh				
1	RMB=0	Device-type Modifier=0						
2	ISO=0		ECMA=0			ANSI=2		
3	RSVD = 0				RDF=2			
4	Additional Length = 0							

Figure 38. INQUIRY DATA Invalid LUN Specified

- **Qualifier** is set to 011b. This indicates the LUN specified in the Command Block is not present.¹
- **Peripheral Dev. Type** is set to 1Fh.
- **Removal Media Bit (RMB)** is always set to zero to indicate no removal media exist.
- **Device-Type Modifier** is set to zero.
- **ISO** is set to zero to indicate that this product does not claim compliance to the International Organization for Standardization (ISO) version of SCSI (ISO DP 9316).
- **ECMA** is set to zero to indicate that this product does not claim compliance to the European Computer Manufacturers Association (ECMA) version of SCSI (ECMA-111).
- **ANSI** indicates the level of the ANSI standard that is supported by the product. The drive supports ANSI SCSI version 2.
- **RDF** is set to two to indicate that the Inquiry Data Format as specified in ANSI SCSI version 2 is supported by the drive.

¹ For all commands, except for inquiry, if an invalid LUN is specified Check Condition status is returned.

7.5.2 Inquiry Data Format when EVPD = 0

Figure 39 shows the data format.

BYTE	BIT							
	7	6	5	4	3	2	1	0
0	Qualifier = 0			Peripheral Device Type = 0				
1	RMB=0	Device-Type Modifier = 0						
2	ISO = 0		ECMA = 0			ANSI = 2		
3	RSVD = 0				RDF = 2			
4	Additional Length = 159 (0x9F)							
5-6	Reserved = 0							
7	REL_A = 0	Wb_32 = 0	Wb_16 = 0	Sync = 1	Link = 1	TTD = 0	CmdQu = 1	SftRe = 0
8-15	Vendor ID = 'IBM'							(ASCII)
16-31	Product ID							(ASCII)
32-35	Product Revision Level							(ASCII)
36-43	Unit Serial Number							(ASCII)
44-55	RAM Microcode P/N							(ASCII)
56-95	Reserved = 0							
96-97	Reserved = ' '							(ASCII)
98-101	Plant of Manufacture							(ASCII)
102-105	Date of Manufacture							(ASCII)
106-107	Reserved = ' '							(ASCII)
108-111	FRU Count = '0001'							(ASCII)
112-113	FRU Field Length = '16'							(ASCII)
114-125	Assembly P/N							(ASCII)
126-135	Assembly EC Level							(ASCII)
136-163	Reserved = 0							

Figure 39. INQUIRY Data - EVPD = 0

- **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 a Direct-Access.
- **Removal Media Bit (RMB)** is always set to zero to indicate no removal media exist.

- **Device-Type Modifier** is set to zero.
- **ISO** is set to zero to indicate that this product does not claim compliance to the International Organization for Standardization (ISO) version of SCSI (ISO DP 9316).
- **ECMA** is set to zero to indicate that this product does not claim compliance to the European Computer Manufacturers Association (ECMA) version of SCSI (ECMA-111).
- **ANSI** indicates the level of the ANSI standard that is supported by the drive. The drive supports ANSI SCSI version 2.
- **RDF** is set to two to indicate that the Inquiry Data Format as specified in ANSI SCSI version 2 is supported by the drive.
- **Additional Length** indicates the number of bytes of inquiry information that follows.
- **REL_A** is set to zero to indicate that the drive does not support 'Relative Address Mode'.
- **Wb_32** is set to zero to indicate that the drive does not support 32-bit wide data transfers.
- **Wb_16** is set to zero to indicate that the drive does not support 16-bit wide data transfers.
- **Sync** is set to one to indicate that the drive supports synchronous data transfer.
- **Link** is set to one to indicate that the drive supports linked commands.
- **TTD** is set to zero to indicate that the drive does not support the CONTINUE I/O PROCESS and TARGET TRANSFER DISABLE message for this logical unit.
- **CmdQu** is set to one to indicate that the drive supports command queuing
- **SftRe** is set to zero to indicate that the drive supports Hard Reset only.
- **Vendor ID** is 'IBM' padded with ASCII blanks.
- **Product ID** is specified in ASCII character.

Product ID	DALS-3540
Assembly P/N	85G3871
Physical Layout	
Bytes per Sector	512
Number of heads	2
Number of disks	1
Maximum LBA	1056767
Number of Blocks	1056768

Figure 40. Product ID vs. Formatted Capacity

- **Product Revision Level** indicates the level of microcode. It indicates ROM microcode level before the media is available and RAM microcode after available.
- **Unit Serial Number** contains the file serial number. If the media is not available, this field will contain ASCII spaces (20h).
- **RAM Microcode P/N** is specified in ASCII characters.
- **Plant of Manufacture** is the plant code of manufacture. If the media is not available, this field will contain ASCII spaces (20h).

Plant of Manufacture	Plant name
0933	Fujisawa
0902	Havant

- **Date of Manufacture** contains 2-digit month code followed by 2-digit year code. If the media is not available, this field will contain ASCII spaces (20h).
- **FRU Count** gives the number of individual Field Replaceable Units within this device. This is the number of entries in the following list where these units are identified by part number.
- **FRU Field length** gives the length in bytes of the list of replaceable units that follows this field.
- **Assembly P/N** gives the drive part number.
- **Assembly EC level** gives the overall drive EC level.
- **FRU P/N** is an internal P/N for the product.

7.5.3 Inquiry Data Format when EVPD = 1 (Page Code = 00h)

BYTE	BIT							
	7	6	5	4	3	2	1	0
0	Qualifier = 0			Peripheral Dev Type = 0				
1	Page Code = 00h							
2	Reserved = 0							
3	Page Length = 04h							
4	First Supported Page Code = 01h							
5	Second Supported Page Code = 03h							
6	Third Supported Page Code = 80h							
7	Fourth Supported Page Code = 82h							

Figure 41. INQUIRY DATA when EVPD = 1 (Page Code = 00h)

- **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 a 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.
- **The Supported Page Code** field contains the Page Codes supported by the target. The list is ascending order.

7.5.4 Inquiry Data Format when EVPD = 1 (Page Code = 01h)

BYTE	BIT							
	7	6	5	4	3	2	1	0
0	Qualifier = 0			Peripheral Dev Type = 0				
1	Page Code = 01h							
2	Reserved = 0							
3	Page Length = 47 (0x2F)							
4	ASCII Length = 24 (0x18)							
5–16	Assembly P/N (ASCII)							
17	0							
18–27	Assembly EC (ASCII)							
28	0							
29–40	Assembly P/N (EBCDIC)							
41–50	Assembly EC (EBCDIC)							

Figure 42. INQUIRY DATA when EVPD = 1 (Page Code = 01h)

- **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 a Direct-Access.
- **Page Code** is set to 1, and this field contains the same value as in the page code field of the INQUIRY command descriptor block.
- **Page length** is set to 47, and this field specifies the length of the following page data.
- **ASCII Length** gives the number of bytes of ASCII data to follow.
- **Assembly P/N (ASCII)** gives, in ASCII, the Part Number for the drive assembly. The field is left aligned and unused bytes are filled with 20h. See section 7.5.2, “Inquiry Data Format when EVPD = 0” on page 52 for actual data.
- **Assembly EC (ASCII)** gives the Engineering Control level for the drive assembly. The field is left aligned and unused bytes are filled with 20h. The data in this field will be updated as design and manufacturing changes are made to the product.
- **Assembly P/N (EBCDIC)** gives, in EBSDIC, the Part Number for the drive assembly. The field is left aligned and unused bytes are filled with 40h. See section 7.5.2, “Inquiry Data Format when EVPD = 0” on page 52 for actual data.
- **Assembly EC (EBCDIC)** gives the Engineering Control level for the drive assembly. The field is left aligned and unused bytes are filled with 40h. The data in this field will be updated as design and manufacturing changes are made to the product.

7.5.5 Inquiry Data Format when EVPD = 1 (Page Code = 03h)

BYTE	BIT							
	7	6	5	4	3	2	1	0
0	Qualifier = 0			Peripheral Dev Type = 0				
1	Page Code = 03h							
2	Reserved = 0							
3	Page Length = 36 (0x24)							
4-7	Reserved = ' ' (ASCII)							
8-11	LOAD ID (ASCII)							
12-15	Mod Level (ASCII)							
16-19	PTF Number = 0							
20-23	Patch Number = 0							
24-35	ROM Code P/N (ASCII)							
36-39	Reserved = 0							

Figure 43. INQUIRY DATA - EVPD = 1 (Page Code = 03h)

- **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 a Direct-Access.
- **Page Code** is set to 3, and this field contains the same value as in the page code field of the INQUIRY command descriptor block.
- **Page length** is set to 36, and this field specifies the length of the following page data.
- **LOAD ID** field is used to verify that the RAM code to be downloaded is compatible with the level of ROM. That is, this field specifies current ROM revision.
- **Mod Level** field specifies the latest RAM code's revision.
- **PTF Number** field is not used and is zero.
- **Patch Number** field is not used and is zero.
- **ROM Code P/N**, in ASCII, the ROM code part number. It is left aligned and unused fields are filled with 20h.

7.5.6 Inquiry Data Format when EVPD = 1 (Page Code = 80h)

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

Figure 44. INQUIRY DATA - EVPD = 1 (Page Code = 80h)

- **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 a Direct-Access.
- **Page Code** is set to 80h, and this field contains the same value as in the page code field of the INQUIRY command descriptor block.
- **Page length** is set to 16, and this field specifies the length of the following page data.
- **Serial Number** gives the drive serial number.

7.5.7 Inquiry Data Format when EVPD = 1 (Page Code = 82h)

BYTE	BIT							
	7	6	5	4	3	2	1	0
0	Qualifier = 0			Peripheral Dev Type = 0				
1	Page Code = 82h							
2	Reserved = 0							
3	Page Length = 51 (33h)							
4	ASCII Length = 28 (1Ch)							
5-9	Product Type (ASCII)							
10	0							
11-15	Model Number (ASCII)							
16	0							
17-25	Serial Number (ASCII)							
26	0							
27-31	Vendor IBM 'IBM' (ASCII)							
32	0							
33-36	Product Type (EBCDIC)							
37-39	Model Number (EBCDIC)							
40	Reserved = 0							
41-48	Serial Number (EBCDIC)							
49-54	Vendor IBM 'IBM' (EBCDIC)							

Figure 45. INQUIRY DATA - EVPD = 1 (Page Code = 82h)

- **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 a Direct-Access.
- **Page Code** is set to 82h, and this field contains the same value as in the page code field of the INQUIRY command descriptor block.
- **Page length** field specifies the length of the following page data.
- **ASCII Length** gives the number of bytes of ASCII data to follow.
- **Product Type (ASCII)** gives 4-digit product type code for the drive. The field is left aligned and unused bytes are filled with 20h.
- **Model Number (ASCII)** gives 3-digit product type code for the drive. The field is left aligned and unused bytes are filled with 20h.

- **Serial Number (ASCII)** gives a serial number of the drive. The field is left aligned and unused bytes are filled with 20h.
- **Vendor ID (ASCII)** gives the vendor as IBM. The field is left aligned and unused bytes are filled with 20h.
- **Product Type (EBCDIC)** gives a 4 digit product type code for the drive. The field is left aligned and unused bytes are filled with 40h.
- **Model Number (EBCDIC)** gives 3-digit product type code for the drive. The field is left aligned and unused bytes are filled with 40h.
- **Serial Number (EBCDIC)** gives the file serial number. The field is left aligned and unused bytes are filled with 40h.
- **Vendor ID (EBCDIC)** gives the vendor as IBM. The field is left aligned and unused bytes are filled with 40h.

7.6 LOG SELECT (4Ch)

LOG SELECT is not supported by the drive.

7.7 LOG SENSE (4Dh)

LOG SENSE is not supported by the drive.

7.8 MODE SENSE (1Ah)

	BIT							
	7	6	5	4	3	2	1	0
BYTE 0	Command code = 1Ah							
BYTE 1	LUN			RSVD = 0				
BYTE 2	PCF		Page Code					
BYTE 3	RSVD = 0							
BYTE 4	Allocation Length							
BYTE 5	VU = 0		RSVD = 0				FLAG	LINK

Figure 46. MODE SENSE (1Ah)

The MODE SENSE command provides a means for the drive to report various device parameters to the initiator. It is the complement to the MODE SELECT command.

Allocation Length indicates the maximum number of bytes 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, then that portion of the data up to the allocation length will be sent. It is noted that this may result in only a portion of a multi-byte field being sent.

7.8.1.1.1 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. The parameters set in the last successful MODE SELECT command.
2. The saved values if a MODE SELECT command has not been executed since the last power-on, hard RESET condition, or BUS DEVICE RESET message .

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 shall be returned containing information that indicate 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 2h, 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 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.

7.8.1.1.2 Page Code: This field specifies which page or pages to return. Page code usage is defined in Figure 47.

Page Code	Description
01h – 38h	Return specific page.
3Fh	Return all available pages.

Figure 47. Page Code Usage

7.8.2 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.

7.8.2.1 Header

	7	6	5	4	3	2	1	0
	BIT							
BYTE 0	Mode Data Length							
BYTE 1	Medium Type = 0							
BYTE 2	WP	Parameter List Length						
BYTE 3	Block Descriptor Length (=0 or 8)							

Figure 48. MODE Parameter List (Header)

- **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 file. (Default medium type)
- **WP.** When used with the MODE SELECT command, the write protected (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.
- **Block Descriptor Length.** The field specifies the length in bytes of the block descriptors.
When used with the MODE SELECT command, zero or eight are supported by the file.
When used with the MODE SENSE command, the file returns eight to indicate that only a single block descriptor is available.

7.8.2.2 Block Descriptor

	7	6	5	4	3	2	1	0
BYTE 0	Density code = 0							
BYTE 1	Number of Blocks (MSB)							
BYTE 2								
BYTE 3	(LSB)							
BYTE 4	RSVD = 0							
BYTE 5	Block Length							
BYTE 6								
BYTE 7								

Figure 49. MODE Parameter Block Descriptor

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 all available blocks
- The exact number of blocks in the data area of the file, which can be obtained with the MODE SENSE command.

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 exact number of blocks.

- **Block Length**

When used with the MODE SELECT command, the **Block length** field must contain 512 or zero, or the file will terminate the command with CHECK CONDITION status.

When used with the MODE SENSE command, the field will return always 512.

7.8.2.3 Page Descriptor

	BIT							
	7	6	5	4	3	2	1	0
BYTE 0	PS	RSVD=0	Density code = 0					
BYTE 1	Page Length							
BYTE 2-n	Mode Parameters							

Figure 50. MODE Parameter Page Descriptor

Each mode page contains a page code, a page length, and a set of mode parameters.

When using the MODE SENSE command, a parameter saveable (PS) bit of one indicates that the mode page can be saved by the file in the reserved area of the file.

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).

The pages supported are described in 7.10, “Mode select data” on page 69.

7.9 MODE SELECT (15h)

	BIT							
	7	6	5	4	3	2	1	0
BYTE 0	Command Code = 15h							
BYTE 1	LUN		PF=1	RSVD = 0			SP	
BYTE 2	RSVD = 0							
BYTE 3	RSVD = 0							
BYTE 4	Parameter List Length							
BYTE 5	VU = 0		RSVD = 0			FLAG	LINK	

Figure 51. MODE SELECT (15h)

The MODE SELECT command provides a means for the initiator to change LUN or device parameters shown below. For further details, please refer to the MODE SENSE command.

- Vendor unique parameters (Unit attention)
- Read/write Error Recovery parameters
- Disconnect/Reconnect control parameters
- Format Device parameters
- Rigid Disk Geometry parameters
- Verify Error Recovery parameters
- Caching parameters
- Control Mode Page
- Power condition parameters

There is only one set of mode select parameters shared by all initiators.

PF A PF (Page Format) bit value of 1 indicates the data sent by the initiator after the mode Mode Select Header and the Block Descriptor, if any, complies to the Page Format. The drive ignores this field since it only accepts mode parameters in the Page Format.

SP Save Page. This indicates;

- 0 b The drive does not save the pages sent during the Data Out phase, however they are used for all following commands until the drive is turned off, reset or another MODE SELECT command is received.
- 1 b The drive saves the data in the reserved area of the disk media. The saved data is used for following commands until another MODE SELECT command is received. The saved parameters are kept and active even after the drive is turned off or reset.

Parameter List Length This specifies the number of bytes to be sent from the initiator. A parameter list length of zero indicates that no data is transferred and is not considered as an error.

The MODE SELECT parameter list contains a four-byte header, followed by zero or one block descriptor, and zero or more pages as shown in the subsequent sections.

7.9.1.1 Application Note

The initiator should issue a MODE SENSE command requesting all Changeable values (see PC field in byte two of the CDB in 7.8, “MODE SENSE (1Ah)” on page 62) prior to issuing a MODE SELECT command. This is necessary to find out which pages and how many length are implemented by the drive. The drive returns the number of bytes supported for each Page in the Pages of the MODE SENSE command. The Page Length set by the initiator in the MODE SELECT command must be the exact value as that returned by the drive in MODE SENSE Page Length. If this is not true, the drive returns Check Condition status with Sense Key of ILLEGAL REQUEST. See 7.8, “MODE SENSE (1Ah)” on page 62.

Note: If an initiator sends MODE SELECT command that changes any parameters applicable to other initiators, the drive generates Unit Attention Condition for all initiators except for the one that issued the MODE SELECT command. The drive sets Additional Sense Code to PARAMETERS CHANGED (2Ah).

7.10 Mode select data

The drive supports the following mode page code:

Page	Description	PS
00h	Vendor Unique Parameters (Unit attention)	1
01h	Read-Write Error Recovery Parameters	1
02h	Disconnect/Reconnect Control Parameters	1
03h	Format Device Parameters	0
04h	Rigid Disk Geometry Parameters	0
07h	Verify Error Recovery Parameters	1
08h	Caching Parameters	1
0Ah	Control Mode Page	1
0Dh	Power Condition Parameters	1

Figure 52. Page Code Usage

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.

7.10.1 Page 0h (Vendor Unique Parameters)

		BIT							
		7	6	5	4	3	2	1	0
BYTE	0	PS	RSVD=0	Page Code = 00h					
BYTE	1	Page Length = 0Eh							
BYTE	2	RSVD=0	UQE	DWD	UAI	RSVD = 0			
BYTE	3	ASDPE	RSVD=0	CMDAC	RPFAE	RSVD = 0			CPE
BYTE	4	Ignore							
BYTE	5	RSVD=0	TCC	DSN	FRDD	DPSDP	RSVD=0	CAEN	LITF
BYTE	6	RSVD = 0							
BYTE	7	RSVD = 0							
BYTE	8	RSVD=0	ADC	QEMC	DRD	LED			
BYTE	9	RSVD = 0							
BYTE	10	Command Aging Limit (Hi byte)							
BYTE	11	Command Aging Limit (Low byte)							
BYTE	12	RSVD = 0							
BYTE	13	RSVD = 0							
BYTE	14	DRRT	DNR	RSVD = 0					
BYTE	15	RSVD = 0							

Figure 53. Page 0h

Fields marked in the table as 'Ignore' are not used or checked by the drive. They will be initialized to zero but may be set as desired. This is for compatibility with older drives.

Changeable Parameter	Default Value
UQE	1
DWD	0
UAI	0
ASDPE	0
CMDAC	0
RPFAE	0
CPE	1
TCC	0

DSN	0
FRDD	1
DPSDP	0
CAEN	1
LITF	0
ADC	1
QEMC	0
DRD	1
LED	0
Command Aging Limit	48
DRRT	0
DNR	1

- **UQE**, The UQE(untagged queuing enable) bit controls whether or not untagged queuing is allowed.
- **DWD**, Disable Write Disconnect bit, is set to zero to indicate that the drive is allowed to disconnect from the SCSI bus after receiving a Write(6), Write extend(10), Write and verify or Write long command and prior to starting the data out phase. (The previous identify message must grant the drive the privilege of disconnection or the drive does not disconnect.)

A DWD bit of one indicates that the drive is not allowed to disconnect from the SCSI bus after receiving a Write(6), Write extend(10), Write and verify or Write long command and prior to the data out phase is started. After the data out phase is started, the drive may disconnect to free the SCSI bus. This occurs if the Target's internal control algorithms and other disconnect/reconnect control parameters indicate that this is appropriate and permissible.

Within the drive, this bit is actually ignored. If no other commands are in progress then the drive will not disconnect. If other commands are in progress then the drive will disconnect to avoid possible dead-lock conditions.

- **UAI**, Unit Attention Inhibit bit. This bit is not used during normal operation while the UAI jumper is removed from the drive. It may however be changed by the user with no effects. If the UAI jumper is added to the drive then this bit controls the generation of unit attention conditions.
- **ASDPE**, Additional Save Data Pointer Enable bit, is used to control the sending of additional save data pointers messages. When set it will cause a save data pointers message to be sent on every disconnection. When not set a save data pointers message is only sent if the current connection contained a data phase and a further data phase will be required to complete the command. This field must be set to zero as this function is not currently supported
- **CMDAC** determines if an LED on the drive is activated while commands are active. As the drive does not support an LED this bit is allowed to be changed by the initiator for host system device driver compatibility.
- **RPF AE**, Report Predictive Failure Analysis Error bit. This bit when set allows the drive to report the result of predictive failure analysis. As this drive does not support predictive failure analysis this bit is ignored. It is allowed to be modified by the initiator for host system device driver compatibility.
- **CPE**, Concurrent processing enable bit, is set to zero to indicate that only untagged and unlinked Request Sense or Inquiry can be executed concurrently.
A CPE bit of one indicates that Read(6), Read extend(10), Write(6) and Write extend(10), as well as the above 2 commands, can be executed concurrently.

- **TCC** Thermal Compensation bit controls when the drive will perform thermal compensation. As this drive does not perform thermal compensation, this bit is ignored internally. It is allowed to be modified by the initiator for host system device driver compatibility.
- **DSN** Disable Target Initiated Synchronous Negotiation is not supported by the drive.
- **FRDD** When set this bit stops the drive reporting a format degraded condition after a failed format or reassign command. As this drive does not have a format degraded mode this bit is ignored. It is allowed to be modified by the initiator for host system device driver compatibility. This bit is ignored internally.
- **DPSDP**, Data Phase Save Data Pointer bit, controls whether the drive sends a save data pointer message at the end of a data phase. When set the drive will send the message after every data phase. When reset the message will only be sent if another data phase is required to complete the command. This bit is ignored internally.
- **CAEN** 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 re-ordered to be executed in on a first come first served basis. When this bit is reset commands are always executed based on the queue re-ordering rules.
- **LITF** This bit disables the drive idle time function which saves the log select counters. If this bit is set these counters are not written to disk and so will be lost at the next power cycle.
- **ADC** Adaptive Cache Enable, when set, allows the drive to modify the read-ahead caching algorithm, ignoring parameters in Page 8h. The adaption is based on analyzing the most recent command history and the current contents of the cache buffers. This bit is ignored internally. The drive will always adapt the read-ahead caching algorithm as if this bit is set.
- **QEMC**, Queue Error Management Control, bit effects how the queue is managed after an error has occurred. If the Qerr bit in page 0Ah is zero then this bit is ignored. If however Qerr is set then this bit controls the action taken with queued commands which originated from initiators other than that which has experienced the error. If this bit is zero then there is no effect on commands from other initiators. If this set commands from all initiators will be aborted. This bit must be set to zero as this function is not currently supported.
- **DRD** Disable Read Disconnect bit, when set, disables the drive from dis-connecting from the bus on a read or read extended command until the contents of the cache have been checked. This bit is ignored internally by the drive. When a read command arrives the drive will automatically disconnect if any other command is in progress. If no commands are in progress then the drive will stay connected until the cache has been checked. It is allowed to be modified by the initiator for host system device driver compatibility.
- **LED** This field is designed to control the operation of a drive LED driver. As the drive does not support an LED the field is internally ignored. It is allowed to be modified by the initiator for host system device driver compatibility.
- **Command Aging Limit** This value is used to control the maximum time a command should wait in the command queue when the CAEN bit is set. Each unit of this timer is 50ms.
- **DRRT** Disable Read Reassign Target, when set, will stop the drive attempting to move the data in a sector to the new location, when the LBA is being reassigned. This bit is internally ignored by the drive which will always attempt to move the data. It is allowed to be modified by the initiator for host system device driver compatibility.
- **DNR** Disable Nested Reassigns, when set, disables nested reassigns and means only the target LBA will be moved on a reassign command. This drive never performs nested reassigns, as it is not required by the disk format, so this bit is ignored internally. It is allowed to be modified by the initiator for host system device driver compatibility.

7.10.2 Page 1h (Read/Write Error Recovery Parameters)

	BIT								
	7	6	5	4	3	2	1	0	
BYTE 0	PS	RSVD=0	Page Code = 01h						
BYTE 1	Page Length = 0Ah								
BYTE 2	AWRE	ARRE	TB	RC = 0	EER=0	PER	DTE	DCR	
BYTE 3	Read Retry Count								
BYTE 4	Correction Span								
BYTE 5	Head Offset Count (Not used)								
BYTE 6	Data Strobe Offset Count (Not used)								
BYTE 7	Reserved								
BYTE 8	Write Retry Count								
BYTE 9	Reserved								
BYTE 10	(MSB)	Recovery Time Limit (Not Used)							
BYTE 11								(LSB)	

Figure 54. Page 1h

Changeable Parameter	Default Value
AWRE	1
ARRE	1
TB	0
PER	0
DTE	0
DCR	0
Correction Span	0
Read Retry Count	01h
Write Retry Count	01h

The Read-Write recovery parameters that will be used during any command that performs a read or write operation to the medium.

- **AWRE**, an Automatic write reallocation enabled bit is set to zero to indicate that the drive shall not perform automatic reallocation of defective data blocks during write operations. A AWRE bit is set to one to indicate that the drive shall perform automatic reallocation of defective data blocks during write operations.

- **ARRE**, an Automatic read reallocation enabled bit is set to zero to indicate that the drive shall not perform automatic reallocation of defective data blocks during read operations.
A ARRE bit is set to one to indicate that the drive shall perform automatic reallocation of defective data blocks during read operations.
- **TB**, Transfer Block bit, is set to one to indicate that a data block that is not recovered within the recovery limits specified shall be transferred to the initiator before CHECK CONDITION status is returned.
A TB bit of zero indicates that such a data block shall not be transferred to the initiator. Data blocks that can be recovered within the recovery limits are always transferred, regardless of the value of the bit.
- **RC**, A read continuous bit. **Must be set to zero** , indicating that the error recovery operations that cause delays are acceptable during the data transfer. Data shall not be fabricated.
- **EER**,An enable early recovery bit. **Must be set to zero** , indicating that the drive shall use an error recovery procedure that minimizes the risk of mis-detection or mis-correction during the data transfer. Data shall not be fabricated.
- **PER**, Post Error bit, is set to one to indicate that the drive reports recovered errors.
- **DTE**, Disable Transfer on Error bit, is set to one to indicate that the drive terminates the DATA phase upon detection of a recovered error .
- **DCR**, Disable Correction bit, is set to one to indicate that Error Correction Code is not used for data error recovery.
A DCR bit of zero indicates that ECC is applied to recover the data.
- **Read Retry Count** sets a limit on the amount of data recovery procedure(DRP) passes the Target attempts when recovering read errors. One pass through DRP involves executing all steps of DRP. Only values of 00h and 01h are valid. A value of zero disables all error recovery procedures.
- **Correction Span** field specifies the size, in bits, of the largest data error burst for which data error correction may be attempted. Any value may be set into this field, including zero. The drive will always use it's default correction capabilities.
- **Head Offset Count** is not supported by the drive.
Note: Head Offset is implemented in the read error recovery routine. The user can not modify the offset value.
- **Write Retry Count** sets a limit on the amount of data recovery procedure(DRP) passes the Target attempts when recovering write errors. One pass through DRP involves executing all steps of DRP. Only values of 00h and 01h are valid. A value of zero disables all error recovery procedures.

The following summarizes valid modes of operation. If an illegal mode is set the mode select command will complete successfully but the action of the drive when an error occurs is undefined.

PER DTE DCR TB DESCRIPTION

0	0	0	0	Retries and Error Correction are attempted. Recovered and/or corrected data (if any) is transferred with no CHECK CONDITION status at the end of the transfer. no err The transfer length is exhausted. soft err The transfer length is exhausted. Transferred data includes blocks containing recovered errors. hard err Data transfer stops when an unrecoverable error is encountered. The unrecoverable block is not transferred to the initiator. The drive then creates the CHECK CONDITION status with the appropriate Sense Key.
0	0	0	1	Retries and Error Correction are attempted. Recovered and/or corrected data (if any) is transferred with no CHECK CONDITION status at the end of the transfer. no err The transfer length is exhausted. soft err The transfer length is exhausted. Transferred data includes blocks containing recovered errors. hard err Data transfer stops when an unrecoverable error is encountered. The unrecoverable block is transferred to the initiator. The drive then creates the CHECK CONDITION status with the appropriate Sense Key.
0	0	1	0	Retries are attempted but no error correction (ECC) is applied. Recovered data (if any) is transferred with no CHECK CONDITION status at the end of the transfer. no err The transfer length is exhausted. soft err The transfer length is exhausted. Transferred data includes blocks containing recovered errors. hard err Data transfer stops when an unrecoverable error is encountered. The unrecoverable block is not transferred to the initiator. The drive then creates the CHECK CONDITION status with the appropriate Sense Key.
0	0	1	1	Retries are attempted but no error correction (ECC) is applied. Recovered data (if any) is transferred with no CHECK CONDITION status at the end of the transfer. no err The transfer length is exhausted. soft err The transfer length is exhausted. Transferred data includes blocks containing recovered errors. hard err Data transfer stops when an unrecoverable error is encountered. The unrecoverable block is transferred to the initiator. The drive then creates the CHECK CONDITION status with the appropriate Sense Key.
0	1	0	0	Illegal Request-DTE must be zero when PER is zero
0	1	0	1	Illegal Request-DTE must be zero when PER is zero
0	1	1	0	Illegal Request-DTE must be zero when PER is zero
0	1	1	1	Illegal Request-DTE must be zero when PER is zero

1	0	0	0	<p>The highest level error is reported at the end of transfer. Retries and error correction are attempted. Recovered and/or corrected data (if any) is transferred with CHECK CONDITION status and RECOVERED ERROR Sense Key set at the end of the transfer.</p> <p>no err The transfer length is exhausted.</p> <p>soft err The transfer length is exhausted. Transferred data includes blocks containing recovered errors. The information byte in the sense data will contain the logical block address of the last recovered error.</p> <p>hard err Data transfer stops when an unrecoverable error is encountered. The unrecoverable block is not transferred to the initiator. The drive then creates the CHECK CONDITION status with the appropriate Sense Key.</p>
1	0	0	1	<p>The highest level error is reported at the end of transfer. Retries and error correction are attempted. Recovered and/or corrected data (if any) is transferred with CHECK CONDITION status and RECOVERED ERROR Sense Key set at the end of the transfer.</p> <p>no err The transfer length is exhausted.</p> <p>soft err The transfer length is exhausted. Transferred data includes blocks containing recovered errors. The information byte in the sense data will contain the logical block address of the last recovered error.</p> <p>hard err Data transfer stops when an unrecoverable error is encountered. The unrecoverable block is transferred to the initiator. The drive then creates the CHECK CONDITION status with the appropriate Sense Key.</p>
1	0	1	0	<p>The highest level error is reported at the end of transfer. Retries are attempted but ECC is not applied. Recovered and/or corrected data (if any) is transferred with CHECK CONDITION status and RECOVERED ERROR Sense Key set at the end of the transfer.</p> <p>no err The transfer length is exhausted.</p> <p>soft err The transfer length is exhausted. Transferred data includes blocks containing recovered errors. The information byte in the sense data will contain the LBA of the last recovered error.</p> <p>hard err Data transfer stops when an unrecoverable error is encountered. The unrecoverable block is not transferred to the initiator. The drive then creates the CHECK CONDITION status with the appropriate Sense Key.</p>
1	0	1	1	<p>The highest level error is reported at the end of transfer. Retries are attempted but ECC is not applied. Recovered and/or corrected data (if any) is transferred with CHECK CONDITION status and RECOVERED ERROR Sense Key set at the end of the transfer.</p> <p>no err The transfer length is exhausted.</p> <p>soft err The transfer length is exhausted. Transferred data includes blocks containing recovered errors. The information byte in the sense data will contain the LBA of the last recovered error.</p> <p>hard err Data transfer stops when an unrecoverable error is encountered. The unrecoverable block is transferred to the initiator. The drive then creates the CHECK CONDITION status with the appropriate Sense Key.</p>

- | | | | | |
|---|---|---|---|---|
| 1 | 1 | 0 | 0 | <p>The highest level error is reported at the end of transfer. Retries and error correction are attempted. Recovered and/or corrected data (if any) is transferred with CHECK CONDITION status and RECOVERED ERROR Sense Key set at the end of the transfer.</p> <p>no err The transfer length is exhausted.</p> <p>soft err The transfer stops on the first soft error detected. The information in the sense data shall contain the LBA of the block in error.</p> <p>hard err Data transfer stops on the unrecoverable error. The drive then creates the CHECK CONDITION status with the appropriate Sense Key.</p> |
| 1 | 1 | 0 | 1 | <p>The highest level error is reported at the end of transfer. Retries and error correction are attempted. Recovered and/or corrected data (if any) is transferred with CHECK CONDITION status and RECOVERED ERROR Sense Key set at the end of the transfer.</p> <p>no err The transfer length is exhausted.</p> <p>soft err The transfer stops on the first soft error detected. The information in the sense data shall contain the LBA of the block in error.</p> <p>hard err Data transfer stops on the unrecoverable error. The unrecoverable error block is returned to the initiator. The drive then creates the CHECK CONDITION status with the appropriate Sense Key.</p> |
| 1 | 1 | 1 | 0 | <p>The highest level error is reported at the end of transfer. Retries are attempted but ECC is not applied. Recovered data is transferred with CHECK CONDITION status and RECOVERED ERROR Sense Key set at the end of the transfer.</p> <p>no err The transfer length is exhausted.</p> <p>soft err The transfer stops on the first soft error detected. The recovered error block is returned to the initiator. The information in the sense data shall contain the logical block address of the block in error.</p> <p>hard err Data transfer stops on the unrecoverable error. The drive then creates the CHECK CONDITION status with the appropriate Sense Key.</p> |
| 1 | 1 | 1 | 1 | <p>The highest level error is reported at the end of transfer. Retries are attempted but ECC is not applied. Recovered and/or corrected data (if any) is transferred with CHECK CONDITION status and RECOVERED ERROR Sense Key set at the end of the transfer.</p> <p>no err The transfer length is exhausted.</p> <p>soft err The transfer stops on the first soft error detected. The information in the sense data shall contain the logical block address of the block in error.</p> <p>hard err Data transfer stops on the unrecoverable error. The unrecoverable error block is returned to the initiator. The drive then creates the CHECK CONDITION status with the appropriate Sense Key.</p> |

7.10.3 Page 2h (Disconnect/Reconnect Parameters)

		BIT								
		7	6	5	4	3	2	1	0	
BYTE	0	PS	RSVD=0	Page Code = 02h						
BYTE	1	Page Length = 0Ah								
BYTE	2	Read Buffer Full Ratio								
BYTE	3	Write Buffer Empty Ratio								
BYTE	4 - 11	Reserved=0								

Figure 55. Page 2h

Changeable Parameter	Default Value
----------------------	---------------

Read Buffer Full Ratio	00h
-------------------------------	-----

Write Buffer Empty Ratio	00h
---------------------------------	-----

The disconnect / reconnect page provides the initiator the means to tune the performance of the SCSI bus.

An initiator may use the IDENTIFY message to grant the drive the general privilege of disconnecting. (Disconnect requests may still be selectively rejected by the initiator by issuing a MESSAGE REJECT).

The drive uses the disconnect/reconnect parameters to control reconnection during READ (operation code 08h and 28h) and WRITE (0Ah , 2Ah and 2E).

- **Read Buffer Full Ratio** is the numerator of a fraction whose denominator is 256. The fraction indicates how full the drive data buffer should be before attempting to reconnect to the SCSI bus. If the ratio is set to 0h, the target will calculate and use an optimal ratio based on the negotiated transfer rate.
- **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 reconnect to the SCSI bus. If the ratio is set to 0h, the target will calculate and use an optimal ratio based on the negotiated transfer rate.

7.10.3.1 Reconnection to a disconnected read command

For a read command, the reconnect is delayed relative to the availability of the first block in the drive data buffer by the fraction of the drive data buffer size. If the remaining data transfer length is less than the fraction of the drive data buffer size, the drive control program calculates the optimal reconnection point to complete the data transfer as early as possible while minimizing the time connected to the SCSI bus.

7.10.3.2 Reconnection to a disconnected write command

For a write command, the Write Buffer Ratio is significant only if the total data transfer length is greater than the size of the drive data buffer. The fraction determines how empty the drive data buffer should be before reconnecting to begin filling the buffer again.

7.10.4 Page 3h (Format Device Parameters)

		BIT							
		7	6	5	4	3	2	1	0
BYTE 0	PS = 0	RSVD=0	Page Code = 03h						
BYTE 1	Page Length = 16h								
BYTE 2	(MSB)	Track per Zone = 0							(LSB)
BYTE 3									
BYTE 4	(MSB)	Alternate Sectors per Zone = 0							(LSB)
BYTE 5									
BYTE 6	(MSB)	Alternate Tracks per Zone = 0							(LSB)
BYTE 7									
BYTE 8	(MSB)	Alternate Tracks per Logical Unit = 0							(LSB)
BYTE 9									
BYTE 10	(MSB)	Sectors per Track = 69h							(LSB)
BYTE 11									
BYTE 12	(MSB)	Data Bytes per Physical Sector = 200h							(LSB)
BYTE 13									
BYTE 14	(MSB)	Interleave = 1							(LSB)
BYTE 15									
BYTE 16	(MSB)	Track Skew Factor = 1Ah							(LSB)
BYTE 17									
BYTE 18	(MSB)	Cylinder Skew Factor = 22h							(LSB)
BYTE 19									
BYTE 20	SSEC	HSEC	RMB	SURF	RESERVED				
BYTE 21-23	RESERVED								

Figure 56. Page 3h

The format device page contains parameters which specify the medium format.

SSEC Zero. Indicates that the drive does not support soft sector formatting.

HSEC One. Indicates that the drive support hard sector formatting.

RMB Zero. Indicates that the media does not support removable. Fixed Disk.

SURF Zero. Indicates that progressive address are assigned to all logical blocks a cylinder prior to allocating address within the next cylinder.

7.10.5 Page 4h (Rigid Disk Drive Geometry Parameters)

	7	6	5	4	3	2	1	0
BYTE 0	RSVD = 0		Page Code = 04h					
BYTE 1	Page Length = 16h							
BYTE 2	(MSB) Number of Cylinders = 1322h							
BYTE 4	(LSB)							
BYTE 5	Number of Heads = 2							
BYTE 6	(MSB) Starting Cylinder-Write Precompensation = 0							
BYTE 8	(LSB)							
BYTE 9	(MSB) Starting Cylinder-Reduced Write Current = 0							
BYTE 11	(LSB)							
BYTE 12	(MSB) Drive Step Rate (Not used)							
BYTE 13	(LSB)							
BYTE 14	(MSB) Landing Zone Cylinder (Not used)							
BYTE 16	(LSB)							
BYTE 17	RESERVED						RPL = 0	
BYTE 18	Rotational Offset = 00 (Not used)							
BYTE 19	RESERVED							
BYTE 20	(MSB) Medium Rotation Rate = 1194h							
BYTE 21	(LSB)							
BYTE 22	RESERVED							
BYTE 23	RESERVED							

Figure 57. Page 4h

The rigid disk drive geometric page specifies various parameters for the drive.

RPL Zero. Indicates that the drive does not support spindle synchronization.

7.10.6 Page 7h (Verify Error Recovery Parameters)

		BIT							
		7	6	5	4	3	2	1	0
BYTE	0	PS	RSVD=0	Page Code = 07h					
BYTE	1	Page Length = 0Ah							
BYTE	2	Reserved = 0			EER=0	PER	DTE=0	DCR	
BYTE	3	Verify Retry Count							
BYTE	4	Correction Span = 00h							
BYTE	5	Reserved = 0							
BYTE	6	Reserved = 0							
BYTE	7	Reserved = 0							
BYTE	8	Reserved = 0							
BYTE	9	Reserved = 0							
BYTE	10	Verify Recovery Time Limit (Not Used)							
BYTE	11								

Figure 58. Page 7h

Changeable Parameter	Default Value
PER	0
DCR	0
Verify Retry Count	01h

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.

Since bytes 4-11 are not changeable, the Mode Select Commands accepts only the values indicated for bytes 4 - 11.

- **EER**, This bit is 0 since the Target does not support early recovery.
- **PER**, See below for description of bit values.
- **DTE**, This bit is 0 since the Target always continues on recovered verify operation errors.
- **DCR**, See below for description of bit values.

PER, DTE, and DCR bit settings in page 7 override those of page 1 during Verify and the Verify portion of Write and Verify. There are only four valid conditions for the PER, DTE, and DCR bits. All other combinations return Check Condition Status.

PER DTE DCR			DESCRIPTION
0	0	0	Soft errors are not reported. ECC is applied to recover the data.
1	0	0	Soft errors are reported. ECC is applied to recover the data.
0	0	1	Soft errors are not reported. ECC is not used to recover the data.
1	0	1	Soft errors are reported. ECC is not used to recover the data.

- **Verify Retry Count** sets a limit on the amount of verify recovery procedure(VRP) passes the Target attempts when recovering verify errors. The Verify Retry Count of one causes the Target to attempt up to one VRP pass per command when a medium error occurs during a verify operation. Only values of 0h and 01h are valid. The value of 0h disables all recovery.
- **Verify Correction Span** field specifies the size, in bits, of the largest data error burst for which data error correction may be attempted. The field may be set to any value but the drive will not use offline correction during verify operations.

7.10.7 Page 8h (Caching Parameters)

		BIT								
		7	6	5	4	3	2	1	0	
BYTE	0	PS	RSVD=0	Page Code = 08h						
BYTE	1	Page Length = 0Ch								
BYTE	2	RESERVED = 0					WCE	MF	RCD	
BYTE	3	Read Retention Priority=0				Write Retention Priority=0				
BYTE	4-5	Disable Pre-fetch Transfer Length								
BYTE	6-7	Minimum Pre-fetch								
BYTE	8-9	Maximum Pre-fetch								
BYTE	10-11	Maximum Pre-fetch Ceiling								
BYTE	12	RESERVED = 0								
BYTE	13	Number of Cache Segments								

Figure 59. Page 8h

Changeable Parameter	Default Value
WCE	1
MF	0
RCD	0
Disable Pre-fetch transfer length	0
Minimum Pre-fetch	0
Maximum Pre-fetch	0
Maximum Pre-fetch Ceiling	0
Number of Cache Segments	4

The caching parameters page defines parameters that affect the use of the cache.

- **WCE**, Write Cache enable bit, is set to zero to indicate that the drive must issue Good Status for Write(6) or Write extend(10) command only after successfully writing the data to the media. A WCE bit of one indicates that the drive may issue Good Status for a Write(6) or Write extend(10) command after successfully receiving the data but before writing it to the media.

Note: When WCE = 1, a Synchronize Cache command must be done to assume data is written to the media before powering down the Target.

- **MF**, Multiplication Factor determines how the Maximum Pre-fetch field is interpreted. When this bit is set the data to pre-fetch is given by the command length multiplied by the value in the Maximum Pre-

fetch field. When this bit is reset the value in the Maximum pre-fetch field is used as the absolute length to pre-fetch. This bit is ignored internally because Adaptive Cache is always enabled.

- **RCD**, Read Cache disable bit, of zero indicates that the drive may return data requested by a READ command by accessing either the cache or the Read Ahead Buffer, or media. A RCD bit of one indicates that the drive shall transfer all data requested by a READ command by accessing the media (i.e., data cannot be transferred from the cache or Read Ahead Buffer).
- **Read Retention Priority**, Demand Read Retention Priority is not supported.
- **Write Retention Priority**, Write Retention Priority is not supported.
- **Disable Pre-fetch Transfer Length** specifies a number of LBA's which if a read command length exceeds will cause the drive not to perform read ahead buffering after the command has completed. A value of zero specifies read ahead should always be performed. This field is ignored internally because Adaptive Cache is always enabled.
- **Minimum Pre-fetch** specifies the minimum number of LBA's that the drive should read ahead after each read command. A value of zero indicates that read ahead should be terminated immediately a new command arrives, except in the case when the new command is on the current head and track. In addition pre-fetch will be abandoned at the end of the current track. This field is ignored internally because Adaptive Cache is always enabled.
- **Maximum Pre-fetch** specifies the maximum number of LBA's to read ahead after a read command. This field can either be used as an absolute value, if the MF bit is 0, or else it will be multiplied by the read command length to give the actual length to read ahead. A value of zero indicates that read ahead will be performed until the segment is full of read ahead data. If the length of a read command is less than the segment size then read ahead will be continued until the segment is full regardless of this setting (except when the read ahead is pre-empted by a subsequent command). This field is ignored internally because Adaptive Cache is always enabled.
- **Maximum Pre-fetch ceiling** specifies the maximum number of blocks the drive should attempt to read ahead. It is particularly relevant when the MF bit is set. A value of zero indicates no limit. If the length of a read command is less than the segment size then read ahead will be continued until the segment is full regardless of this setting (except when the read ahead is pre-empted by a subsequent command). This field is ignored internally because Adaptive Cache is always enabled.
- **Number of Cache Segments**

This field is used to indicate to the drive how many segments are requested by the initiator. This field is ignored internally. Number of segments are always 4. Segments size are 16KB.

7.10.8 Page Ah (Control Mode Page Parameters)

		BIT								
		7	6	5	4	3	2	1	0	
BYTE	0	PS	RSVD=0	Page Code = 0Ah						
BYTE	1	Page Length = 6								
BYTE	2	RESERVED = 0							RLEC = 0	
BYTE	3	Queue Algorithm Modifier			RESERVED = 0		QErr	DQue		
BYTE	4	EECA = 0	RESERVED = 0			RAENP = 0	UAAENP = 0	EAENP = 0		
BYTE	5	RESERVED = 0								
BYTE	6	Ready AEN Holdoff Period = 0								
BYTE	7									

Figure 60. Page Ah

Changeable Parameter	Default Value
Queue Algorithm Modifier	0
QErr	00h
DQue	0

Following are parameter options for Page 0A of MODE SELECT.

- **RELC**, Report Log Exception Condition, bit of zero specifies that the target shall not report log exception conditions. This bit must be zero.
- **Queue algorithm modifier** specifies restrictions on the algorithm used for re-ordering commands that are tagged with the SIMPLE QUEUE TAG message.
 - 0h : Restricted re-ordering. The target shall re-order the actual execution sequence of the queued commands from each initiator such that data integrity is maintained for that initiator.
 - 1h : Un-restricted re-ordering allowed. The target may re-order 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 re-ordering is disabled
 - 9-Fh : RESERVED
- **QErr**, Queue Error Management, bit of zero specifies that the Target suspends execution of queued and active commands from any Initiator which receives a Check Condition Status until pending sense data is cleared. Those commands still queued after the Target has returned Check Condition Status, continue execution in a normal manner when the pending status is cleared. A QErr bit of one specifies that all active commands and all queued commands from all initiators are aborted when the Target returns the Check Condition Status. A unit attention condition will be generated for each initiator which had commands in the queue except the initiator that received the Check Condition Status. The sense key will be

set to Unit Attention and the additional sense code will be set to **COMMANDS CLEARED BY ANOTHER INITIATOR**.

- **DQue**, Disable Queuing, bit of zero specifies that tagged queuing shall be enabled if the target supports tagged queuing. A DQue bit of one specifies that tagged queuing shall be disabled. Any queue commands for that I_T_L nexus shall be aborted. Any subsequent queue tag message received shall be rejected with a MESSAGE REJECT message and I/O process shall be executed as an untagged command.
- **EECA**, Enable Extended Contingent Allegiance, bit must be zero.
- **RAENP**, Ready AEN Permission, must be zero.
- **UAAENP**, Unit Attention AEN Permission, must be zero.
- **EAENP**, Error AEN Permission, must be zero.
- **Ready AEN Holdoff Period** must be zero.

7.10.9 Page 0Dh (Power Condition)

		BIT								
		7	6	5	4	3	2	1	0	
BYTE 0	PS	RSVD=0	Page Code = 0Dh							
BYTE 1	Page Length = 0Ah									
BYTE 2	Reserved = 0									
BYTE 3	Reserved = 0						RSVD = 0	Standby		
BYTE 4	Reserved = 0									
BYTE 5										
BYTE 6										
BYTE 7										
BYTE 8	(MSB)	Standby Condition Timer								
BYTE 9										
BYTE 10										
BYTE 11	(LSB)									

Figure 61. Page 0Dh

The power condition page provides the initiator the means to control the length of time a target will delay before changing its power requirements. There is no notification to the initiator that a target has entered into one of the power conditions. On the receipt of a command, only the timer(s) controlling power for the command shall be reset and then restarted on completion of that command.

- **Standby.** A Standby bit of one indicates a target shall use the Standby Condition Timer to determine the length of inactivity time to wait before entering the Standby condition. A standby bit of zero indicates a target shall not enter the Standby condition.
- **The Standby Condition Timer.** The Standby condition Timer field indicates the inactivity time in 100 millisecond increments that the target shall wait before entering the Standby condition.

7.11 PRE-FETCH (34h)

	7	6	5	4	3	2	1	0
BYTE 0	Command Code = 34h							
BYTE 1	LUN			Reserved = 0			Immed	RelAdr = 0
BYTE 2	(MSB) Logical Block Address (LSB)							
BYTE 3								
BYTE 4								
BYTE 5								
BYTE 6	Reserved = 0							
BYTE 7	(MSB) Transfer Length (LSB)							
BYTE 8								
BYTE 9	VU = 0	Reserved = 0				FLAG	LINK	

Figure 62. Pre-Fetch (34h)

The PRE-FETCH command requests the drive to transfer data to the cache. No data is transferred to the initiator.

Immed

Immediate.

When the Immediate (Immed) bit of the CDB is zero and:

- If an error occurs while reading, error recovery procedures are attempted. The drive returns GOOD status or CHECK CONDITION status based on the set parameters of the MODE SELECT Page 1.
- If there is enough room in the segment for all of the Requested Data or if the Transfer Length is zero and no error occurs while reading, the drive returns CONDITION MET status when the command completes.
- If there is not enough room in the segment and the Transfer Length is not zero and no error occurred while reading, the drive returns GOOD status when the command completes.

When the Immediate (Immed) bit of the CDB is one and:

- If there is enough room in the segment for all of the Requested Data or if the Transfer Length is zero, the drive returns CONDITION MET status as soon as the CDB is verified.
- If there is not enough room in the segment and the Transfer Length is not zero, the drive returns GOOD status as soon as the CDB is verified.
- The reading of data is handled like Read-Ahead operation. This implies the PRE-FETCH is terminated upon receipt of another command.
- if an error is encountered:
 1. The drive terminates the PRE-FETCH operation and does not attempt to recover the data.

2. The error is not reported to the Initiator for the current command. The error is reported during the next command if the next command is a Read command and requests the block which encountered the error.)
3. The blocks which were successfully read prior to the errored block are retained in the cache.

RelAdr

Relative Block Address. **Must be set to zero** to indicate that the logical block address field specifies the first logical block of the range of logical blocks to be operated by this command. Relative address is not supported.

Transfer length

The transfer length field specifies the number of contiguous blocks of data that are to be transferred into the cache. A transfer length of zero indicates that logical blocks are to be transferred into the cache until the segment is filled or there are no more blocks on the media.

7.12 READ (08h)

	BIT							
	7	6	5	4	3	2	1	0
BYTE 0	Command Code = 08h							
BYTE 1	LUN			(MSB)	LBA			
BYTE 2	LOGICAL BLOCK ADDRESS							
BYTE 3	LOGICAL BLOCK ADDRESS (LSB)							
BYTE 4	TRANSFER LENGTH							
BYTE 5	VU = 0	RSVD = 0			FLAG	LINK		

Figure 63. READ (08h)

The READ command requests the drive to transfer the specified number of blocks of data to the initiator starting at the specified logical block address.

Logical block address This field specifies the logical block address from which the read operation begins.

Transfer length This field specifies the number of blocks to be transferred. A value of zero implies that 256 blocks (512 bytes/block) are to be transferred.

Note: Errors are handled by ERP (Error Recovery Procedure). ERP is controlled by the error recovery parameters specified by MODE SELECT command.

7.13 READ CAPACITY (25h)

	7	6	5	4	3	2	1	0
BYTE 0	Command Code = 25h							
BYTE 1	LUN			RSVD = 0			RelAdr	
BYTE 2	(MSB) Logical Block Address							
BYTE 3								
BYTE 4								
BYTE 5	(LSB)							
BYTE 6	RSVD = 0							
BYTE 7	RSVD = 0							
BYTE 8	RSVD = 0							PMI
BYTE 9	VU = 0		RSVD = 0			FLAG	LINK	

Figure 64. READ CAPACITY (25h)

The READ CAPACITY command returns information regarding a capacity of the drive.

- **RelAdr.** A Relative Address is not supported and must be set to zero.
- **Logical Block Address** is used in conjunction with the PMI bit.
- **PMI**, Partial Medium Indicator indicates;
 - 0 The drive returns the last logical block address.
 - 1 The drive returns the last logical block address on the same track where the LBA in the command is located. This option provides the information initiators need to determine the amount of space available on the same track which is accessible without a head switch or seek.

7.13.1.1 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.

	7	6	5	4	3	2	1	0
BYTE 0	(MSB) Logical Block Address							
BYTE 1								
BYTE 2								
BYTE 3	(LSB)							
BYTE 4	(MSB) Block Length = 200h							
BYTE 5								
BYTE 6	= 512							
BYTE 7	(LSB)							

Figure 65. Format of READ CAPACITY command reply

- **Block Length** specifies the length in bytes of the block. It is set to 200h (512).

7.14 READ DEFECT DATA (37h)

	7	6	5	4	BIT		3	2	1	0
BYTE 0	COMMAND CODE = 37h									
BYTE 1	LUN			Rsvd = 0					0	
BYTE 2	Rsvd = 0			Plist	Glist	Defect List Format				
BYTE 3	Rsvd = 0									
BYTE 4										
BYTE 5										
BYTE 6										
BYTE 7	Allocation length (MSB)									
BYTE 8	(LSB)									
BYTE 9	VU = 0		RSVD = 0					FLAG	LINK	

Figure 66. Read Defect Data (37h)

The READ DEFECT DATA command requests that the drive transfers the medium defect data to initiators.

If the drive is unable to access any medium defect data, it returns Check Condition status with the appropriate sense key. The sense key is set to either MEDIUM ERROR(03h) if a medium error occurred or NO SENSE (00h) if the list does not exist with the additional sense code set to DEFECT LIST ERROR (19h).

Plist The Primary Defect List (Plist) bit set to one indicates that the drive returns the primary list of defects. A Plist bit of zero indicates that the drive does not return the Primary Defect list.

Glist The Grown Defect List (Glist) bit set to one indicates that the drive returns the Grown Defect List. A Glist bit of zero indicates that the drive does not return the Grown Defect List of defects.

Note: When both bits are set to one, both Plist and Glist are returned. When both bits are set to zero, the drive returns a four-byte Defect List Header only.

Defect List Format The Defect List Format Field is used by initiators to indicate the preferred format for the defect list.

The Defect List Format of '100b (Bytes from Index Format) ' and '101b (Physical Sector Format)' are supported. If the requested format is not supported by the drive, it returns the defect list in its default format '101' and terminates the command with Check Condition status. The sense key is set to Recovered Error(01h) and the Additional Sense Code is set to Defect List Not Found(1Ch).

The drive sends Defect Descriptors after a four-byte Defect List Header.

The drive transfers all of the Read Defect Data up to the number of bytes allocated by the initiator.

Note: The drive terminates 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 four-byte header, followed by zero or more defect descriptors.

7.14.1 Defect List Header

	7 6 5 4 3 2 1 0 BIT Defect List Header				
BYTE 0	Rsvd = 0				
BYTE 1	<table border="1" style="width: 100%; border-collapse: collapse;"> <tr> <td style="width: 25%;">Rsvd = 0</td> <td style="width: 10%;">Plist</td> <td style="width: 10%;">Glist</td> <td style="width: 55%;">Defect List Format</td> </tr> </table>	Rsvd = 0	Plist	Glist	Defect List Format
Rsvd = 0	Plist	Glist	Defect List Format		
BYTE 2	Defect List Length (MSB)				
BYTE 3	(LSB)				

Figure 67. Defect List Header

7.14.2 Bytes from Index Format (100b)

	Defect Descriptors
BYTE 0	(MSB)
BYTE 1	Cylinder Number of Defect
BYTE 2	(LSB)
BYTE 3	Head Number of Defect
BYTE 4	(MSB)
BYTE 5	Defect Bytes from Index
BYTE 6	
BYTE 7	(LSB)

Figure 68. Defect Descriptors of Bytes from Index Format

Defect Bytes from Index is gotten using the following equation:

$$\text{Defect Bytes from Index} = (\text{Physical Sector Numbers}) \times N$$

Where: N = Bytes per sector (512 bytes)

7.14.3 Physical Sector Format (101b)

Defect Descriptors	
BYTE 0	Cylinder Number of Defect (MSB)
BYTE 1	
BYTE 2	
BYTE 3	Head Number of Defect
BYTE 4	Defective Sector Number (MSB)
BYTE 5	
BYTE 6	
BYTE 7	
8-byte sets of defect list follow.	

Figure 69. Defect Descriptor of Physical Sector Format

The Defect List Length field of Defect List Header 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.

If the Allocation Length is insufficient to transfer all of the defect descriptors, the Defect List Length is not adjusted to reflect the truncation and the drive returns CHECK CONDITION status.

7.15 READ EXTENDED (28h)

	7	6	5	4	3	2	1	0
BYTE 0	Command Code = 28h							
BYTE 1	LUN			DPO = 0	FUA	Reserved = 0		RelAdr = 0
BYTE 2	(MSB) Logical Block Address							
BYTE 3								
BYTE 4								
BYTE 5	(LSB)							
BYTE 6	Reserved = 0							
BYTE 7	(MSB) Transfer Length							
BYTE 8	(LSB)							
BYTE 9	VU = 0	Reserved = 0				FLAG	LINK	

Figure 70. Read Extended (28h)

The READ EXTENDED command requests the drive to transfer data to initiators. The Transfer Length field specifies the number of contiguous logical blocks of data that is transferred.

- DPO** Disable Page Out. **must be set to zero** Disable Page Out is not supported.
- FUA** Force Unit Access. An FUA bit of 1 indicates that the data is read from the media and not from the cache. An FUA bit of 0 allows the data to be read from either the media or the cache.
- RelAdr** Relative Block Address. **must be set to zero**, to indicate that the logical block address field specifies the first logical block of logical blocks to be operated by this command. Relative Address is not supported.
- Transfer length** The number of contiguous blocks to be transferred. If the transfer length is zero, the seek is made but no data is transferred. This condition is not considered an error. If read ahead is enabled, read ahead is started after the seek completes.

7.16 READ BUFFER (3Ch)

	7	6	5	4	3	2	1	0
BYTE 0	Command Code = 3Ch							
BYTE 1	LUN			RSVD = 0		MODE		
BYTE 2	Buffer ID = 0							
BYTE 3	(MSB) Buffer Offset							
BYTE 4								
BYTE 5	(LSB)							
BYTE 6	(MSB) Allocation length							
BYTE 7								
BYTE 8	(LSB)							
BYTE 9	VU = 0		RSVD = 0			FLAG		LINK

Figure 71. READ BUFFER (3Ch)

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.

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

000 Read combined header and data

010 Read Data

011 Read Buffer Descriptor

All others Not supported.

7.16.1 Combined Header And Data (Mode 000b)

In this mode, a four-byte header followed by data bytes are returned to initiators 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 an initiator, whichever is less.

The four-byte READ BUFFER header (Figure 72 on page 99) is followed by data bytes from the data buffer of the drive.

					BIT				
	7	6	5	4	3	2	1	0	
BYTE 0	RSVD = 0								
BYTE 1	(MSB)	Buffer Capacity							
BYTE 2									
BYTE 3								(LSB)	

Figure 72. READ BUFFER Header

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 by WRITE BUFFER command.

Following the READ BUFFER header, the drive transfers data from its data buffer.

7.16.2 Read Data (Mode 010b)

In this mode, the DATA IN phase contains buffer data.

Buffer ID	This field must be set to zero. If other value is specified, CHECK CONDITION status is returned with Sense Key to ILLEGAL REQUEST and Additional Sense Code to ILLEGAL FIELD IN CDB.
Buffer Offset	This 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, CHECK CONDITION status is returned with 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.

7.16.3 Descriptor (Mode 011b)

In this mode, a maximum of four bytes of READ BUFFER descriptor are returned. The drive returns the descriptor information for the buffer specified by the buffer ID. (See 'READ DATA' mode for the buffer ID.)

Buffer ID	If there is no buffer associated with the specified buffer ID, the drive returns all zeros in the READ BUFFER descriptor.
Buffer Offset	This field is reserved.
Allocation Length	This must be set to four or greater. The drive transfers the lesser of the allocation length or four bytes of READ BUFFER descriptor. The READ BUFFER descriptor is defined in Figure 73 on page 100.

	7	6	5	4	3	2	1	0
	BIT							
BYTE 0	Offset Boundary							
BYTE 1	Buffer Capacity							
BYTE 2	(MSB)							
BYTE 3	(LSB)							

Figure 73. READ BUFFER DESCRIPTOR

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, i.e. sector boundary of 512 bytes.

7.17 READ LONG (3Eh)

	BIT							
	7	6	5	4	3	2	1	0
BYTE 0	Command Code = 3Eh							
BYTE 1	LUN		Reserved = 0		CORT = 0	RelAdr = 0		
BYTE 2	(MSB)							
BYTE 3	LOGICAL BLOCK ADDRESS							
BYTE 4								
BYTE 5	(LSB)							
BYTE 6	Reserved							
BYTE 7	(MSB)							
BYTE 8	Byte Transfer Length (LSB)							
BYTE 9	VU = 0	RSVD = 0			FLAG	LINK		

Figure 74. READ LONG (3Eh)

The READ LONG command requests the drive to transfer **one block** of data to the initiator. The transfer data includes;

- 512 bytes of data
- ECC field data
- **CORT**
 - 0 A corrected bit of zero causes the logical block to be read without any Error Correction Code (ECC) made by the drive.
 - 1 Not supported. A corrected bit of one causes the data to be corrected by ECC before transferring the data to the initiator.
- **RelAdr** Relative Block Address is not supported.
- **LOGICAL BLOCK ADDRESS** field specifies the logical block at which the read operation is made.
- **Byte Transfer Length** field must exactly specify 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 drive terminates the command with CHECK CONDITION status, Sense key set to ILLEGAL REQUEST and Additional Sense Code set to INVALID FIELD IN CDB.

The transfer length is calculated as follows:

$$\text{Transfer length} = \text{logical block size} + 16 = 528$$

7.18 REASSIGN BLOCKS (07h)

	BIT							
	7	6	5	4	3	2	1	0
BYTE 0	Command Code = 07h							
BYTE 1	LUN			RSVD = 0				
BYTE 2	RSVD = 0							
BYTE 3	RSVD = 0							
BYTE 4	RSVD = 0							
BYTE 5	VU = 0		RSVD = 0			FLAG	LINK	

Figure 75. REASSIGN BLOCKS (07h)

The REASSIGN BLOCKS command requests the drive to reassign a logical block to an available spare. The REASSIGN BLOCKS command attempts to allocate spare blocks. The logical block address is transferred to the drive during the DATA OUT phase. One to four block(s) may be specified for relocation per REASSIGN BLOCKS command.

Reassignment completes upon the completion of the REASSIGN BLOCKS command. At this time, the defective logical block address has been added to the Grown (“G” list) Defect List.

Data contained at the logical block address being reassigned is not preserved by the drive, and is filled with a constant pattern.

Following is the format of the data sent by the initiator during the DATA OUT phase:

	BIT							
	7	6	5	4	3	2	1	0
BYTE 0	RSVD = 0							
BYTE 1	RSVD = 0							
BYTE 2 BYTE 3	(MSB) Defect list length = 4/8/12/16 (LSB)							
BYTE 4 BYTE 5 BYTE 6 BYTE 7	(MSB) Defective Logical Block Address No.1 (LSB)							
BYTE 8 BYTE 9 BYTE 10 BYTE 11	(MSB) Defective Logical Block Address No.2 (LSB)							
BYTE 12 BYTE 13 BYTE 14 BYTE 15	(MSB) Defective Logical Block Address No.3 (LSB)							
BYTE 16 BYTE 17 BYTE 18 BYTE 19	(MSB) Defective Logical Block Address No.4 (LSB)							

Figure 76. Format of REASSIGN BLOCKS data

Note: If the drive finds a defective block by verifying ECC, the drive does not start the REASSIGN BLOCKS process, and returns CHECK CONDITION status with Sense Key of MEDIUM ERROR.

- **Defect List Length** must be 4,8,12 or 16. Otherwise, the drive returns Check Condition with Sense Key of Illegal request.
- **Defective logical block address** is four-byte length. The initiator can specify from one to four Defective logical blocks. Defective logical block addresses must be ordered in ascending order or the drive returns Check Condition status.

7.19 RECEIVE DIAGNOSTIC (1Ch)

	7	6	5	4	3	2	1	0
BYTE 0	Command Code = 1Ch							
BYTE 1	LUN			RSVD = 0				
BYTE 2	RSVD = 0							
BYTE 3	(MSB) Parameter List Length (LSB)							
BYTE 4								
BYTE 5	VU = 0		RSVD = 0				FLAG	LINK

Figure 77. RECEIVE DIAGNOSTIC (1Ch)

The Receive Diagnostic command requests that analysis data requested by a Send Diagnostics command be sent to the initiator.

- **Parameter List 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.

7.19.1 Receive Diagnostic Page 0

This page contains a list of supported pages.

Byte	7	6	5	4	3	2	1	0
0	Page Code = 0							
1	Reserved = 0							
2 - 3	Page Length = 02h							
4	First supported page 0							
5	Second supported page 40							

7.19.2 Receive Diagnostic Page 40h

translate address. Using the Send

Using the Send Diagnostics 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.

Byte	7	6	5	4	3	2	1	0
0	Page Code = 40h							
1	Reserved = 0							
2 - 3	Page Length = 0Ah							

Byte	7	6	5	4	3	2	1	0
4	Reserved = 0					Supplied Format		
5	RA	ALTS	ALTT	Reserved = 0		Translate Format		
6 - 13	Translated Address							

- **Supplied Format** is the value supplied by the Send Diagnostic command it may be one of the two following values
 - **000b** block format
 - **101b** physical sector format

It specifies the format in which the address has been supplied.
- **Translate Format** is the value supplied by the Send Diagnostic command it may take either of the two formats specified above and specifies the format that the address has been translated into List.
- **RA (Reserved Area)** is set to 1 if the translated block is a reserved area.
- **ALTS (Alternate Sector)** is set to 1 if the translated block is in alternate sector area.
- **ALTT (Alternate Track)** is set to 1 if the translated block is in alternate track area. This bit is not used by the drive.
- **Translated Address** contains the address in the translate format. If it is a logical block address it is contained within the first 4 bytes of the field, i.e. bytes 6 to 9 of the page data. For a physical format it is as follows.

Byte	7	6	5	4	3	2	1	0
6 - 8	Cylinder Number							
9	Head Number							
10 - 13	Sector Number							

7.20 RELEASE (17h)

	7	6	5	4	3	2	1	0
BYTE 0	Command Code = 17h							
BYTE 1	LUN		3rdPty	3rd Party ID			Ext=0	
BYTE 2	Reservation Identification							
BYTE 3	RSVD = 0							
BYTE 4	RSVD = 0							
BYTE 5	VU = 0		RSVD = 0			FLAG	LINK	

Figure 78. RELEASE (17h)

The RELEASE command is used to release a LUN previously reserved.

Note: It is not an error for an initiator to release a LUN that is not currently reserved.

- **3rdPty** bit indicates that :
 - 1 This release process is for a third party which is specified by 3rd Party ID.
 - 0 This release process is for the initiator itself.
- **3rd Party ID** specifies the ID of the third party for which the LUN is reserved.²
- **Extents** must be 0. Extension is not supported.
- **Reservation Identification** field is ignored.

² Refer 7.22, “RESERVE (16h)” on page 108

7.21 REQUEST SENSE (03h)

	BIT							
	7	6	5	4	3	2	1	0
BYTE 0	Command Code = 03h							
BYTE 1	LUN			RSVD = 0				
BYTE 2	RSVD = 0							
BYTE 3	RSVD = 0							
BYTE 4	ALLOCATION LENGTH							
BYTE 5	VU = 0		RSVD = 0			FLAG	LINK	

Figure 79. REQUEST SENSE (03h)

The REQUEST SENSE command requests the drive to transfer sense data of 32 bytes.

The sense data is available on following conditions,

- The previous command to the specified I_T_L nexus terminated with CHECK CONDITION status. I_T_L nexus exists between an initiator, a target (drive) and a logical unit.
- Other information (e.g. medium position) is available in any fields.
- The previous command to the specified I_T_L nexus ended unexpected BUS FREE error.

If REQUEST SENSE command with an invalid LUN is received, the drive returns GOOD status with Sense Key of ILLEGAL REQUEST and Additional Sense Code of LOGICAL UNIT NOT SUPPORTED.

If the drive has no sense data available to return, it returns Sense Key of NO SENSE and Additional Sense Code of NO ADDITIONAL SENSE INFORMATION.

The sense data shall be preserved by the drive for the initiator until retrieved by the REQUEST SENSE command or until any other command for the same I_T_L nexus. Sense data shall be cleared upon receipt subsequent command including REQUEST SENSE to the same I_T_L nexus.

Separate sense data is maintained by the device for each initiator. Therefore there is no requirement for an initiator to expeditiously clear CHECK CONDITION status as this does not affect other initiators in a multi-initiator system.

The drive returns the number of bytes in the allocation length or 32 bytes whichever is less.

7.22 RESERVE (16h)

	7	6	5	4	3	2	1	0
BYTE 0	Command Code = 16h							
BYTE 1	LUN		3rdPty	3rd Party ID			Ext=0	
BYTE 2	Reservation Identification							
BYTE 3	Extent List Length = 0							
BYTE 4	(MSB) (LSB)							
BYTE 5	VU = 0		RSVD = 0			FLAG	LINK	

Figure 80. RESERVE (16h)

The RESERVE command is used to reserve a LUN for an initiator. This reservation can be either for;

- the initiator which sends this command.
- the third party which is specified by this command.

This command results in reserving the entire LUN for the initiator until one of the following occurs:

- The reservation is superseded by another valid RESERVE command from the initiator that made the reservation.
- The LUN is released by a RELEASE command from the same initiator.
- A hard reset condition occurs. (A SCSI bus Reset assertion)
- A BUS DEVICE RESET message is received from any initiator.
- Power off/on occurs.
- **3rdPty** bit is to indicates that :
 - 1** This reservation is for a third party which is specified by 3rd Party ID.
 - 0** This reservation is for the initiator itself.
- **3rd Party ID** specifies the ID of the third party for which the LUN is reserved.

Note: The LUN may be only released by the initiator who sent the RESERVE command.
- **Extents** must be 0. Extension is not supported.
- **Reservation Identification** is ignored.
- **Extent List** length must be zero. Extent List length is ignored.

7.22.1 Command Processing While Reserved

When a reservation for a LUN is made by an initiator for the same or a different initiator (Third Party option), the following rules apply.

- 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 that receive the reservation then :
 - A REQUEST SENSE or INQUIRY command are permitted.
 - A RELEASE command is permitted but is ignored.
 - Any other command result 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 :
 - A REQUEST SENSE, INQUIRY, RELEASE or RESERVE command are permitted.
 - Any other commands result 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 :
 - RESERVE command result in a RESERVATION CONFLICT status.
 - A RELEASE command is permissible but is ignored.
 - Any other command is executed.

Note: A 3rdPty bit has no meaning if the initiator designated itself.

7.23 REZERO UNIT (01h)

	BIT							
	7	6	5	4	3	2	1	0
BYTE 0	Command Code = 01h							
BYTE 1	LUN			RSVD = 0				
BYTE 2	RSVD = 0							
BYTE 3	RSVD = 0							
BYTE 4	RSVD = 0							
BYTE 5	VU = 0		RSVD = 0			FLAG	LINK	

Figure 81. REZERO UNIT (01h)

The REZERO UNIT command requests that the drive seeks to logical block address 0.

7.24 SEEK (0Bh)

	7	6	5	4	3	2	1	0
BYTE 0	Command Code = 0Bh							
BYTE 1	LUN			(MSB)	LBA			
BYTE 2	LOGICAL BLOCK ADDRESS							
BYTE 3	LOGICAL BLOCK ADDRESS (LSB)							
BYTE 4	Reserved							
BYTE 5	VU = 0	RSVD = 0			FLAG	LINK		

Figure 82. SEEK (0Bh)

The SEEK command requests the drive to seek to the specified logical block address.

7.25 SEEK EXTENDED (2Bh)

	7	6	5	4	3	2	1	0
BYTE 0	Command Code = 2Bh							
BYTE 1	LUN			RSVD = 0				0
BYTE 2	(MSB) Logical Block Address							
BYTE 3								
BYTE 4								
BYTE 5	(LSB)							
BYTE 6	RSVD = 0							
BYTE 7	RSVD = 0							
BYTE 8	RSVD = 0							
BYTE 9	VU = 0		RSVD = 0				FLAG	LINK

Figure 83. SEEK EXTENDED (2Bh)

The SEEK EXTENDED command requests the drive to seek to the specified logical block address.

7.26 SEND DIAGNOSTIC (1Dh)

	7	6	5	4	3	2	1	0
BYTE 0	Command Code = 1Dh							
BYTE 1	LUN		PF	RSVD=0	SlfTst	DevOf1	UntOf1	
BYTE 2	RSVD = 0							
BYTE 3	(MSB)		Parameter List Length				(LSB)	
BYTE 4								
BYTE 5	VU = 0		RSVD = 0			FLAG	LINK	

Figure 84. SEND DIAGNOSTIC (1Dh)

The SEND DIAGNOSTIC command requests the file 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 1 indicates the data sent by the Initiator conform to the page structure as specified in SCSI-2 standard. This bit must be set to 1 if the SlfTst bit is set to 0. This bit is ignored by the Target if the SlfTst bit is set.
- **SlfTst** set to 1 indicates the device should perform it's internal self test. If set to 0 then a parameter list should be sent by the initiator.
- **DevOf1** this bit is ignored by the drive for compatibility.
- **UntOf1** this bit is ignored by the drive for compatibility.
- **Parameter List Length** is ignored by the drive if the SlfTst bit is set. Otherwise 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 generated with a Sense Key of Illegal Request and additional sense of Invalid Field in CDB.

If the SlfTst bit is set upon command completion, the following status is returned:

- GOOD status for successful test completion.
- CHECK CONDITION status for unsuccessful test completions.

The self diagnostics consists of two parts :³

- The first part is executed immediately after power up. This test is performed to verify all hardware which is not related to the disk drive. The local microprocessor, RAM (scratchpad and buffer), and control electronics are included here.
- The second part is executed after the spindle motor is started. This includes disk access (seek), R/W channel, and error correction circuitry verification. A reserved area on the disk is used for this test.

³ See 10.19.2, "Diagnostics Command" on page 169 for a detailed listing of the operations carried out by the Diagnostics Command.

Both tests are performed as a result of the SEND DIAGNOSTIC command. The SEND DIAGNOSTICS will fail with CHECK CONDITION status if it is issued while the spindle motor is not turning. (Such as after STOP command has been received.)

Note: The self diagnostic is also performed at Power On Reset time.

Note: The SCSI bus signals will not be corrupted when the device is executing the SEND DIAGNOSTIC command.

7.26.1 Send Diagnostic Page 0

This page requests that the file return a list of supported pages on the next receive diagnostics command.

Byte	7	6	5	4	3	2	1	0
0	Page Code = 0							
1	Reserved = 0							
2 - 3	Page Length = 0							

7.26.2 Send Diagnostic Page 40h

the translate address page. This

This allows the initiator to translate a logical block address 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 Diagnostic command.

Byte	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 either of the two following values
 - **000b** block format
 - **101b** physical sector format

It specifies the format in which the address has been supplied.

- **Translate Format** may take either of the two formats specified above and specifies that format that the address should be translated into. 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 of Illegal Request and Illegal Field in Parameter List.

- **Address to Translate** contains the address to translate. If the logical block format is specified then the first 4 bytes of the field, i.e. bytes 6 to 9, contain the LBA and the remainder must be zero. For the physical format the address must be specified as follows.

Byte	7	6	5	4	3	2	1	0
6 - 8	Cylinder Number							
9	Head Number							
10 - 13	Sector Number							

7.27 START/STOP UNIT (1Bh)

	7	6	5	4	3	2	1	0
BYTE 0	Command Code = 1Bh							
BYTE 1	LUN			RSVD = 0			Immed	
BYTE 2	RSVD = 0							
BYTE 3	RSVD = 0							
BYTE 4	RSVD = 0							Start
BYTE 5	VU = 0		RSVD = 0			FLAG		LINK

Figure 85. START/STOP Unit (1Bh)

The START/STOP UNIT command is used to spin up or stop the spindle motor.

- **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 after a spin-up.
- **Start** bit is to specify:
 - 0** Stop the spindle.
 - 1** Start the spindle.

Note: Once the drive has become ready (after a power on) the Start/Stop UNIT command can be used without any errors, regardless of the state of the motor, i.e. stopped or spinning.

7.28 SYNCHRONIZE CACHE (35h)

	7	6	5	4	3	2	1	0
BYTE 0	Command Code = 35h							
BYTE 1	LUN			RSVD = 0			Immed = 0	RelAdr = 0
BYTE 2	(MSB)							
BYTE 3	Logical Block Address							
BYTE 4								
BYTE 5	(LSB)							
BYTE 6	Reserved							
BYTE 7	(MSB)							
BYTE 8	Number of Blocks							
	(LSB)							
BYTE 9	VU = 0		RSVD = 0				FLAG	LINK

Figure 86. SYNCHRONIZE CACHE (35h)

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 begins.
- **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 are within the range.
- **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 Check Condition status with Sense Key of Illegal Request and Additional Sense Code of Invalid Field in CDB.
- **RelAdr** (relative address) must be zero.
The drive does not support the relative addressing.
If the RelAdr bit is set to one, the drive returns Check Condition status with Sense Key of Illegal Request and Additional Sense Code of Invalid Field in CDB.

7.29 TEST UNIT READY (00h)

	BIT							
	7	6	5	4	3	2	1	0
BYTE 0	Command Code = 00h							
BYTE 1	LUN			RSVD = 0				
BYTE 2	RSVD = 0							
BYTE 3	RSVD = 0							
BYTE 4	RSVD = 0							
BYTE 5	VU = 0		RSVD = 0			FLAG		LINK

Figure 87. TEST UNIT READY (00h)

The TEST UNIT READY command allows the initiator to check if the drive is READY. The SCSI specification defines READY as the condition that the drive accepts media-access commands without returning CHECK CONDITION status.

The drive verifies that the motor is spinning at the specified speed (4500 RPM).

- If the spindle motor is not spinning at the correct speed, CHECK CONDITION status is returned with sense key of NOT READY.
- If the motor is spinning at the specified speed, the drive accepts normal 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 START UNIT command with an Immediate bit of 1. In this mode the START UNIT command returns COMMAND COMPLETE status before the completion of motor spin-up and expects the initiator to issue TEST UNIT READY commands to determine when the motor has reached the proper speed.

Note: The spindle automatically starts in default. The drive does not execute any commands other than TEST UNIT READY, INQUIRY or REQUEST SENSE command until the the Power On sequence is completed. The drive returns CHECK CONDITION status with Sense Key of NOT READY and Sense Code of IN PROCESS OF BECOMING READY for all other commands during the Power On period.

7.30 VERIFY (2Fh)

	7	6	5	4	3	2	1	0
BYTE 0	Command Code = 2Fh							
BYTE 1	LUN			RSVD = 0			ByteChk	0
BYTE 2	(MSB) Logical Block Address							
BYTE 3								
BYTE 4								
BYTE 5	(LSB)							
BYTE 6	RSVD = 0							
BYTE 7	(MSB) Verification Length							
BYTE 8	(LSB)							
BYTE 9	VU = 0		RSVD = 0				FLAG	LINK

Figure 88. VERIFY (2Fh)

The VERIFY command requests the drive to 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** indicates;

- 0** The verification is performed by ECC check. No data transfer from the initiator is performed in this case. If an ECC error is detected, CHECK CONDITION status is returned with sense key of MEDIUM ERROR.
- 1** Byte-by-byte comparison is not supported.

7.31 WRITE (0Ah)

	BIT							
	7	6	5	4	3	2	1	0
BYTE 0	Command Code = 0Ah							
BYTE 1	LUN			(MSB)	LBA			
BYTE 2	LOGICAL BLOCK ADDRESS							
BYTE 3	LOGICAL BLOCK ADDRESS (LSB)							
BYTE 4	TRANSFER LENGTH							
BYTE 5	VU = 0		RSVD = 0			FLAG	LINK	

Figure 89. WRITE (0Ah)

The WRITE command requests the drive to write the specified number of blocks of data from the initiator to the medium starting at the specified logical block address.

See 7.12, “READ (08h)” on page 91 for the parameters.

7.32 WRITE EXTENDED (2Ah)

	7	6	5	4	3	2	1	0
BYTE 0	Command Code = 2Ah							
BYTE 1	LUN			DPO = 0	FUA	RSVD = 0		RelAdr = 0
BYTE 2	(MSB) Logical Block Address							
BYTE 3								
BYTE 4								
BYTE 5	(LSB)							
BYTE 6	RSVD = 0							
BYTE 7	(MSB) Transfer Length							
BYTE 8	(LSB)							
BYTE 9	VU = 0		RSVD = 0			FLAG		LINK

Figure 90. WRITE EXTENDED (2Ah)

The WRITE EXTENDED command requests the drive to write the data transferred from the initiator. This command is processed like the standard WRITE command except for the longer transfer length.

DPO Disable page out. **must be set to zero** Disable page out is not supported.

FUA Force unit access. A FUA bit of 1 indicates that the drive must write the data to the media before returning Good Status. An FUA bit of 0 indicates the drive may return Good Status prior to writing the data to the media.

RelAdr Relative Block Address. **must be set to zero**, indicating that the logical block address field specifies the first logical block of the range of logical blocks to be operated on by this command. Relative address is not supported.

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.

7.33 WRITE AND VERIFY (2Eh)

	7	6	5	4	3	2	1	0
BYTE 0	Command Code = 2Eh							
BYTE 1	LUN			DPO = 0	RSVD = 0		ByteChk = 0	RelAdr = 0
BYTE 2	(MSB) Logical Block Address							
BYTE 3								
BYTE 4								
BYTE 5	(LSB)							
BYTE 6	RSVD = 0							
BYTE 7	(MSB) Transfer Length							
BYTE 8	(LSB)							
BYTE 9	VU = 0		RSVD = 0				FLAG	LINK

Figure 91. WRITE AND VERIFY (2Eh)

WRITE AND VERIFY command requests the drive to write the data transferred from the initiator to the medium and to verify that the data is correctly written.

- **ByteChk** the options:

ByteChk	Description
----------------	--------------------

0	The data is read back from the disk 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 of MEDIUM ERROR.
----------	--

1	Not supported.
----------	----------------

- DPO(Disable page out) **must be set to zero** Disable page out is not supported.
- Relative Block Address **must be set to zero** indicating that the logical block address field specifies the first logical block of logical blocks to be operated by this command. Relative address is not supported.
- A transfer length of zero indicates that no data is transferred.
- If caching is enabled, the command performs an implied Force Unit Access (FUA) and an implied Synchronize Cache before starting the operation. This insures that the medium, not the cache, is being verified.

7.34 WRITE BUFFER (3Bh)

	7	6	5	4	3	2	1	0
BYTE 0	Command Code = 3Bh							
BYTE 1	LUN			RSVD = 0		MODE		
BYTE 2	Buffer ID							
BYTE 3	(MSB) Buffer Offset							
BYTE 4								
BYTE 5	(LSB)							
BYTE 6	(MSB) Parameter list length							
BYTE 7								
BYTE 8	(LSB)							
BYTE 9	VU = 0		RSVD = 0			FLAG		LINK

Figure 92. WRITE BUFFER (3Bh)

The WRITE BUFFER command is used in conjunction with the READ BUFFER command as a diagnostic function for testing the drive memory and the SCSI bus integrity. This command does not alter the medium of the drive. Additional modes are provided for downloading microcode and for downloading 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

- 000** Write combined header and data
- 010** Data
- 100** Download Microcode
- 101** Download Microcode and Save
- All other modes are not supported by the drive.

7.34.1 Combined Header And Data (Mode 000b)

In this mode, the data to be transferred is preceded by a four-byte header.

Buffer ID

This field must be zero. If another value is specified, no download function are performed and 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.

Buffer Offset This field must be zero. If another value is specified, no download function are performed and 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.

Parameter List Length This field 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.

The four-byte header consists of all reserved bytes.

	7	6	5	4	3	2	1	0
BYTE 0	RSVD = 0							
BYTE 1	RSVD = 0							
BYTE 2	RSVD = 0							
BYTE 3	RSVD = 0							

Figure 93. WRITE BUFFER Header

7.34.2 Write Data (Mode 010b)

In this mode, the DATA OUT phase contains buffer data.

Buffer ID This 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 This 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.

Parameter List Length This field specifies the Parameter List Length. It must be;

- less than the capacity of the buffer size.
- on a sector boundary. In other words, it must be a multiple of 512.

If a invalid 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.

7.34.3 Download Microcode (Mode 100b)

In this mode, the microcode is transferred to the control memory space of the file. Once downloaded the drive will operate with the newly downloaded code until the next power cycle.

Buffer ID The buffer ID field is used to indicate which portion of the microcode image is being downloaded. If it is set to '00'x then the main microprocessor code is updated. If the buffer ID is '01'x then the HDC pico code is updated.

Any other value for the buffer ID except '00'x and '01'x will cause the command to terminate with CHECK CONDITION status. The drive shall set

sense key to ILLEGAL REQUEST and additional sense code to ILLEGAL FIELD IN CDB.

Buffer Offset

This specifies the starting address of the downloaded Microcode. It must be zero. If an invalid 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.

Transfer Length

The total microcode length must be 0x8000 for both the main micro processor and the HDC pico processor updates. It may also be set to '0000'x in which case no code is updated. If an invalid 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.

This process generates a unit attention condition for MICROCODE HAS BEEN CHANGED for all initiators except the one which sent the write buffer command. Once the write buffer command has been completed the new microcode is immediately ready for operation.

Note: This option is documented within this spec to indicate that the drive will accept a command of this form though it is not expected that a customer will ever issue such a command. To use the write buffer command in this manner a special microcode version will be 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 operations after downloading.

If the write buffer command is executed with this option and the code downloaded is not valid for the drive then it is to be expected that the drive will hang on this or a subsequent command. This should normally be recoverable by a power on/off cycle.

7.34.4 Download Microcode and Save (Mode 101b)

In this mode, the microcode is transferred to the drive and saved into the System reserved area on the disk. The downloaded code becomes effective immediately after download and after each Power On Reset until it is overwritten by another download microcode and save option.

The code must be downloaded to the drive in 5 separate blocks. Each of these being 32KB in size, giving a total code size of 160KB. The 5 blocks must be sent in the correct sequential order but other commands may be interspersed between them, however if power is cycled or a reset occurs the download procedure must be restarted. As the blocks are received they are stored into a reserved area of the drive - at this stage they DO NOT overwrite the current microcode. When the last of the 5 blocks is received the download will be checked via checksum and signature for integrity and compatibility with the ROM code. Once these checks have passed the drive will overwrite the old code stored on the disk with the new download. The new download will also be loaded into memory.

Buffer ID

The buffer ID field is used to indicate which portion of the microcode image is being downloaded.

The buffer ID must be set to '00'x on the first block, '01'x on the second block etc. The blocks must be sent to the drive in the correct logical order, with no omissions or resends of blocks. On all write buffer commands except the final one the drive merely stores the code and no action is taken. When the final block is received, the drive attempts to use the new microcode. In between downloading the code blocks the drive will continue to accept and process other commands from this and any other initiator. It is not until the final block of code has been received that the code the processors is running is changed. In addition it is legal to send some but not all of the code and then to start sending a different level of code as long as this second level of code is started from the beginning, ie buffer ID '00'x.

Any value for the buffer ID except '00'x or an incrementing value based on the last block sent will cause the command to terminate with CHECK CONDITION status. The drive shall set sense key to ILLEGAL REQUEST and additional sense code to ILLEGAL FIELD IN CDB.

Buffer Offset

This specifies the starting address of the downloaded Microcode. It must be zero. If an invalid 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.

Transfer Length

Total microcode length must be specified. This value should be 0x8000. It may also be set '0000'x in which case no code is updated. If an invalid 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.

Note: New code to be downloaded to the drive will be provided by development either in request to a customer's request for additional function or as a result of a bug fix to a critical drive function. Please note however that not all possible fixes or new function can be applied to the drive in this manner and that there is a very considerable dependency on the level of ROM code contained within the drive. If invalid code or code that is not compatible with the ROM code is downloaded the drive will normally reject this code and will continue normal operation. However there is a very small possibility of invalid code being accepted (about 1 in 4E9) and if this occurs the unit will usually become inoperable and have to be returned to the manufacturer to be recovered.

This process generates a unit attention condition of MICROCODE HAS BEEN CHANGED for all initiators except the one sending the last write buffer command (It is legal to send the different blocks of code from varying initiators as long as they arrive at the drive in the correct logical order - although this is not recommended). Once the final write buffer command has completed the new code is immediately available and will be used for processing following commands. There is no delay required after the completion of the write buffer command before the drive can start accepting new commands. However the final write buffer command may take up to 5 seconds to complete.

7.35 WRITE LONG (3Fh)

	BIT							
	7	6	5	4	3	2	1	0
BYTE 0	Command Code = 3Fh							
BYTE 1	LUN			Reserved = 0			RelAd = 0	
BYTE 2	(MSB)							
BYTE 3	LOGICAL BLOCK ADDRESS							
BYTE 4								
BYTE 5	(LSB)							
BYTE 6	Reserved = 0							
BYTE 7	(MSB)							
BYTE 8	Byte Transfer Length (LSB)							
BYTE 9	VU = 0		RSVD = 0			FLAG		LINK

Figure 94. WRITE LONG (3Fh)

The WRITE LONG command requests the drive to write **one block** of data transferred from the initiator. The transfer data must include;

- 512 bytes of data
- 16 bytes of ECC data

Parameters are;

- **RelAd** (Relative Block Address). This is not supported.
- **LOGICAL BLOCK ADDRESS** field specifies the logical block at which the read operation shall occur.
- **Byte Transfer Length**. This field must exactly specify 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, Check Condition status is returned with Sense Key of ILLEGAL REQUEST and Additional Sense Code of INVALID FIELD IN CDB.

8.0 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.




Figure 95. SCSI Status Byte. Format of the SCSI STATUS byte. All Reserved fields(R) are set to zero.

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 file. The REQUEST SENSE command should be issued to determine the nature of the condition.
- 04h** **CONDITION MET**
This status indicates that the requested operation is satisfied. (See "Pre-Fetch" Command)
- 08h** **BUSY**
This condition is returned when disconnect privilege is not granted while the file 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 reissue the command and grant the disconnect privilege.
- 10h** **INTERMEDIATE/GOOD**
This status is returned for every command in a series of linked commands (except the last command), unless an error, exception, or abnormal condition causes a CHECK CONDITION status or a RESERVATION CONFLICT status to be set. If this status is not returned, the chain of linked commands is broken, and no further commands in the series are executed.
- 14h** **INTERMEDIATE/CONDITION MET**
This status is the combination of CONDITION MET and INTERMEDIATE /GOOD.
- 18h** **RESERVATION CONFLICT**
This status is returned whenever an SCSI device attempts to access the file, but it has been reserved by another initiator. (See 7.22, "RESERVE (16h)" on page 108.)
- 28h** **QUEUE FULL**
This status indicates that the targets command queue is full. If 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 is not valid.

9.0 SCSI MESSAGE SYSTEM

This chapter details how the message system is implemented on the file. Included is a functional description of the supported messages.

9.1 Supported Messages

The message supported by the file is shown in Figure 96.

MESSAGE	CODE (hex)	Direction	Negate ATN Before last ACK	
COMMAND COMPLETE	00	IN	—	
SYNCHRONOUS DATA TRANSFER REQUEST	010301	IN	OUT	Yes
SAVE DATA POINTER	02	IN	—	
RESTORE POINTERS	03	IN	—	
DISCONNECT	04	IN	—	
DISCONNECT	04		OUT	Yes
INITIATOR DETECTED ERROR	05		OUT	Yes
ABORT	06		OUT	Yes
MESSAGE REJECT	07	IN	OUT	Yes
NO OPERATION	08		OUT	Yes
MESSAGE PARITY ERROR	09		OUT	Yes
LINKED COMMAND COMPLETE	0A	IN	—	
LINKED COMMAND COMPLETE (w/FLAG)	0B	IN	—	
BUS DEVICE RESET	0C		OUT	Yes
ABORT TAG	0D		OUT	Yes
CLEAR QUEUE TAG	0E		OUT	Yes
SIMPLE QUEUE TAG	20XX	IN	OUT	No
HEAD OF QUEUE TAG	21XX		OUT	No
ORDERED QUEUE TAG	22XX		OUT	No
IDENTIFY	80-FF	IN	—	
IDENTIFY	80-FF		OUT	No

Key: IN = Target to Initiator, OUT = Initiator to target.
YES = Initiator shall negate ATN before last ACK of message.
NO = Initiator may or may not negate ATN before last ACK of message.
— = Not applicable
XX = Queue Tag

Figure 96. Supported Messages

If an unsupported message is received, the drive will send the *MESSAGE REJECT* message to the initiator. When unsupported messages are received to a valid nexus of the drive, the drive executes a following command. When unsupported messages are received to an invalid nexus of the drive, the drive goes to BUS FREE phase.

9.1.1 COMMAND COMPLETE (00)

The file sends this message to the initiator to indicate that the execution of a command has terminated and that valid status has been sent to the initiator. After successfully sending this message, the file releases all bus signals and goes to BUS FREE phase.

9.1.2 SYNCHRONOUS DATA TRANSFER REQUEST (01,03,01H)

Byte	Value	Description
0	01H	Extended message
1	03H	Extended message length
2	01H	SYNCHRONOUS DATA TRANSFER REQUEST code
3	M	Transfer period (M times 4 nanoseconds)
4	X	REQ/ACK offset

Figure 97. Synchronous Data Transfer Request.

A pair of Synchronous Data Transfer Request (SDTR) messages shown in Figure 97 are exchanged between an initiator and a Target to establish the synchronous data transfer mode between the two devices. The message exchange establishes the permissible transfer period and REQ/ACK offset for a synchronous data transfer between the two devices. The initiator may initiate a synchronous data transfer negotiation at any time after the LUN has been identified. A Synchronous Data Transfer Request(SDTR) message exchange shall be initiated by an SCSI device whenever a previously arranged data transfer agreement may have become invalid. *The agreement becomes invalid after any condition which may leave the data transfer agreement in an indeterminate state such as;*

1. after a Power-on Reset
2. after a SCSI Bus "hard" reset condition
3. after a Bus Device Reset message

In addition, a SCSI device may initiate a SDTR message exchange whenever it is appropriate to negotiate a new data transfer agreement(either synchronous or asynchronous).

M The transfer period(M above) is the minimum time allowed between leading edges of successive REQ pulses and of successive ACK pulses to meet the device requirements for successful reception of data. The file supports transfer period in the range 100 nSec to 475 nSec in 25 nSec increments.

REQ/ACK Offset

The ACK/REQ offset(X above) is the maximum number of REQ pulses allowed to be outstanding before the leading edge of its corresponding ACK pulses is received at the file. A REQ/ACK offset value of zero indicate asynchronous data transfer mode. The file supports REQ/ACK offset values in the range 0 through 15.

If ATN is negated before all bytes of a multiple-byte extended message is received, the file will go to **BUS FREE** to signal a catastrophic error.

9.1.2.1 Synchronous Negotiation Started by the Initiator

The file responds to each Initiator requested transfer period as shown in the following figure Figure 98:

Initiator Request	Target Response	Target Transfer Period	Maximum Burst Rate
0 <= Mi <= 25	Mt = 25	100 nSec	10.00 MB/s
26 <= Mi <= 31	Mt = Mi	125 nSec	8.00 MB/s
32 <= Mi <= 37	Mt = Mi	150 nSec	6.67 MB/s
38 <= Mi <= 43	Mt = Mi	175 nSec	5.71 MB/s
44 <= Mi <= 50	Mt = Mi	200 nSec	5.00 MB/s
51 <= Mi <= 56	Mt = Mi	225 nSec	4.44 MB/s
57 <= Mi <= 62	Mt = Mi	250 nSec	4.00 MB/s
63 <= Mi <= 68	Mt = Mi	275 nSec	3.64 MB/s
69 <= Mi <= 75	Mt = Mi	300 nSec	3.33 MB/s
76 <= Mi <= 81	Mt = Mi	325 nSec	3.08 MB/s
82 <= Mi <= 87	Mt = Mi	350 nSec	2.86 MB/s
88 <= Mi <= 93	Mt = Mi	375 nSec	2.67 MB/s
94 <= Mi <= 100	Mt = Mi	400 nSec	2.50 MB/s
101 <= Mi <= 106	Mt = Mi	425 nSec	2.35 MB/s
107 <= Mi <= 112	Mt = Mi	450 nSec	2.22 MB/s
113 <= Mi <= 118	Mt = Mi	475 nSec	2.11 MB/s
119 <= Mi <= 255	Mt = Mi	(Asynchronous mode)	N/A

Figure 98. Initiator Request/Target Response

9.1.2.2 Synchronous Negotiation Started by the Target

If the file recognize that negotiation is required, the file sends a SDR message to the initiator with transfer period equal to 200 nSec (M = 50). The file interprets the Initiator corresponding transfer period as shown in the following figure Figure 99:

Initiator's Response	Target Transfer Period	Maximum Burst Rate
0 <= Mi <= 24	Send Message Reject(Async mode)	N/A
25 <= Mi <= 25	100 nSec	10.00 MB/s
26 <= Mi <= 31	125 nSec	8.00 MB/s
32 <= Mi <= 37	150 nSec	6.67 MB/s
38 <= Mi <= 43	175 nSec	5.71 MB/s
44 <= Mi <= 50	200 nSec	5.00 MB/s
51 <= Mi <= 56	225 nSec	4.44 MB/s
57 <= Mi <= 62	250 nSec	4.00 MB/s
63 <= Mi <= 68	275 nSec	3.64 MB/s
69 <= Mi <= 75	300 nSec	3.33 MB/s
76 <= Mi <= 81	325 nSec	3.08 MB/s
82 <= Mi <= 87	350 nSec	2.86 MB/s
88 <= Mi <= 93	375 nSec	2.67 MB/s
94 <= Mi <= 100	400 nSec	2.50 MB/s
101 <= Mi <= 106	425 nSec	2.35 MB/s
107 <= Mi <= 112	450 nSec	2.22 MB/s
113 <= Mi <= 118	475 nSec	2.11 MB/s
119 <= Mi <= 255	Send Message Reject(Async mode)	N/A

Figure 99. Target Response to Initiator's Transfer Period

9.1.3 SAVE DATA POINTER (02)

This message is sent from the file to direct the initiator to copy the active data pointer to the saved data pointer. The SAVE DATA POINTER message is only sent if the initiator has previously indicated the ability to accommodate disconnection and reconnection via the IDENTIFY message .

The file will send the SAVE DATA POINTER message to the initiator prior to sending a DISCONNECT message to the initiator if a data phase has occurred and another data phase is required to successfully complete the command.

9.1.4 RESTORE POINTERS (03)

This message is sent from the file to direct an initiator to copy the most recently saved pointers to the corresponding command, data, and status pointers. Command and status pointers should be restored to the beginning of the present command and status areas. The data pointer should be restored to the value at the beginning of the data area in the absence of a SAVE DATA POINTER message or to the value at the point at which the last SAVE DATA POINTER message occurred. Also see 9.4, “SCSI Bus Related Error Handling Protocol” on page 139.

9.1.5 DISCONNECT (04)

This message is sent from the file to inform an initiator that the present connection is going to be broken. A later reconnect will be required in order to complete the current command. The disconnection is to free the SCSI bus while the file performs a relatively long operation that does not require the bus. These messages are only sent if the initiator previously indicated (via the IDENTIFY message) the ability to accommodate disconnection and reconnection.

The DISCONNECT message may also be sent from the initiator to the file to disconnect from the SCSI bus. If the file supports disconnecting at the time the DISCONNECT message is received from the initiator, the file will switch to the MESSAGE IN phase, send a DISCONNECT message to the initiator (possibly preceded by a SAVE DATA POINTER message), and then go to the BUS FREE phase. The file will not participate in another ARBITRATION phase for at least a disconnection delay. If the file does not support disconnecting at the time the Disconnect message is received from the initiator, the file will respond by sending a MESSAGE REJECT message to the initiator.

9.1.6 INITIATOR DETECTED ERROR (05)

This message is sent from an initiator to inform the file that an error has been detected that does not preclude the file from retrying the previous COMMAND, DATA and STATUS phase. The source of the error may be either related to previous activities on the SCSI bus or may be internal to the initiator and unrelated to any previous SCSI bus activity

If the initiator intends to send this message, the initiator must assert the ATN signal prior to its release of ACK for the last byte transferred in the information phase that is to be retried. This provides an interlock so the file can determine which information phase to retry.

After receiving this message, the file may retry the previous phase by sending a RESTORE POINTERS message to the initiator and then repeating the previous COMMAND, DATA, or STATUS phase.

9.1.7 ABORT (06)

This message is sent from the initiator to direct the file to clear the present operation for this initiator and logical unit, including queued command(s). If a logical unit has been identified, then all pending data and status for the issuing initiator and this logical unit will be cleared and the file will go to the BUS FREE phase. Pending data and status for other logical unit and initiators will not be cleared. If a logical unit has not been identified, the file will go to the BUS FREE phase without affecting an operation on any logical unit for this initiator or any other initiator. In either case, no status or ending message will be sent to the initiator for this operation. It is not an error to send the ABORT message to a logical unit that is not currently performing an operation for the initiator.

Note: It is permissible for an initiator to select the file/LUN after the file has disconnected from the initiator, for the purpose of sending an IDENTIFY message followed by an ABORT message. This will abort the command on the specified logical unit.

9.1.8 MESSAGE REJECT (07)

This message is sent from either the initiator or the file to indicate that the last message received was inappropriate or has not been implemented.

If the initiator intends to send this message, the initiator must assert the ATN signal prior to its release of ACK for the REQ/ACK handshake of the message byte that is to be rejected. This provides an interlock so the file can determine which message is rejected.

If the file intends to send this message, the file will change to the MESSAGE IN phase and send the MESSAGE REJECT message to the initiator prior to transferring any additional message bytes (or any other information phase bytes) from the initiator regardless of ATN signal. This provides an interlock so the initiator can determine which message is rejected. After the file sends a MESSAGE REJECT message and if ATN signal is still asserted then it shall return to the MESSAGE OUT phase. the subsequent MESSAGE OUT phase shall begin with first byte of a message.

9.1.9 NO OPERATION (08)

This message is sent from the initiator to the file when the initiator does not currently have any other valid message to send. This message is ignored by the file and will not affect any operation.

9.1.10 MESSAGE PARITY ERROR (09)

This message is sent from the initiator to inform the file that the last message byte received had a parity error.

If the initiator intends to send this message, the initiator must assert the ATN signal prior to its release of ACK for the REQ/ACK handshake of the message byte that has the parity error. This provides an interlock so the file can determine which message byte has the parity error.

If the file receives this message under any other circumstance, the file will change to BUS FREE to signal a catastrophic error. After receiving this message, the file will retry sending the previous message to the initiator.

9.1.11 LINKED COMMAND COMPLETE (0A)

The file sends this message to the initiator to indicate that execution of a linked command (with flag bit equal to zero) has completed and that valid status has been sent to the initiator. After successfully sending this message, the file goes to COMMAND phase to receive the next command.

9.1.12 LINKED COMMAND COMPLETE WITH FLAG (0B)

The file sends this message to the initiator to indicate that the execution of a linked command with flag bit set to one has completed and that valid status has been sent to the initiator. After successfully sending this message, the file goes to COMMAND phase to receive the next command.

9.1.13 BUS DEVICE RESET (0C)

This message is sent from an initiator to direct the file to clear all current commands. This message forces a hard reset condition which will reset the file to an initial state with no operations pending for any initiator. After receiving this message, the file will go to the BUS FREE phase.

9.1.14 ABORT TAG (0D)

When the target receives this message successfully, it clears the current I/O process and go to Bus Free. If the target has already started execution of I/O process, the execution will be halted. Pending status, data and commands for other active or queued I/O process shall not be affected.

9.1.15 CLEAR QUEUE TAG (0E)

All I/O process for all initiators shall be cleared. All active I/O process shall be terminated. The target shall go to the Bus Free phase following successfully receipt of this message.

9.1.16 QUEUE TAG MESSAGES(20h, 21h, 22h)

Byte	Value	Description
0	20H	Simple Queue Tag message
	21H	Head of Queue Tag message
	22H	Ordered Queue Tag message
1	XXh	Queue Tag

Figure 100. Queue Tag Messages

Queue Tag messages are used to specify an identifier, called a Queue Tag, for an I/O process which establish the I_T_L_Q nexus. The queue tag filed is an 8-bit unsigned integer assigned by the initiator during an initial connection. The Queue Tag for every I/O process for each I_T_L nexus must be unique. If the target receives a Queue Tag that is currently in use for the I_T_L nexus it will respond as "Incorrect Initiator Response". A Queue Tag becomes available for re-assignment when I/O process ends. The numeric value of a Queue Tag has no effect on the order of execution.

Whenever an initiator connects to the target, the appropriate Queue Tag message must be sent immediately following the Identify message and within the same MESSAGE OUT phase to establish the I_T_L_Q nexus for the I/O process.

Whenever the target reconnects to an initiator to continue a tagged I/O process, the Simple Queue Tag message is sent immediately following the Identify and within the same MESSAGE IN phase to revive the I_T_L_Q nexus for the I/O process.

9.1.16.1 SIMPLE QUEUE TAG (20h)

The Simple Queue Tag Message specifies that the current I/O process be placed in the command queue. The order of execution, with respect to other I/O processes received with Simple Queue Tag Messages, is up to the discretion of the target. The target will send a Simple Queue Tag Messages after reselection for I/O processes that were received with either Simple, Ordered, or Head of Queue Tag messages.

9.1.16.2 HEAD OF QUEUE TAG (21h)

Commands with this tag should be inserted into the head of the queue. When a command is being executed, this tagged command will be inserted to the head of queue to be executed after the command being currently executed. The previous command being executed will not be terminated by this tagged command. This tagged command will wait until the previous command is through. If plural head-of-queue tagged commands are received, those commands will be executed in LIFO (Last in First out) order.

9.1.16.3 ORDERED QUEUE TAG (22h)

This tagged command is executed in the order received. All commands received before this command should be executed before this command, and all commands received after this command should be executed after this command.

9.1.17 IDENTIFY (80 - FF)

This message is set by either the initiator or the file to establish the logical path connection between the two devices.

The IDENTIFY message is defined as follows:

- Bit 7** This bit is always set to one to distinguish the IDENTIFY message from other messages.
- Bit 6** This bit is only set to one by the initiator to grant the file the privilege of disconnecting. If this bit is zero, the file will not disconnect, unless the initiator instructs the file to disconnect by sending a DISCONNECT Message to the file. This bit is set to zero when the file sends an IDENTIFY message to the initiator.
- Bits 5-3** These bits are reserved and must be zero for an IDENTIFY message.
Note: If an invalid Identify message is received with these bits not equal to zero, then the file sends a MESSAGE REJECT message to the initiator and goes to the Bus Free phase to signal a catastrophic error condition.
- Bits 2-0** These bits specify the logical unit number (LUN).
Only one LUN may be identified for any one selection sequence. If the file receives an IDENTIFY message with a new LUN after the LUN had previously been identified, the file will go to the BUS FREE phase to signal a catastrophic error. The initiator may send more than one Identify message during a selection sequence in order to toggle disconnect/reconnect permission if the specified LUN remains the same.

When the IDENTIFY message is sent from the file to the initiator during reconnection, an implied RESTORE POINTERS message must be performed by the initiator.

9.2 Supported Message Functions

The implementation of the supported messages will also include the following functions.

- Retry SCSI Command, DATA IN, DATA OUT, or STATUS phase
The retry will be caused by the following error condition.
 - The file detected SCSI bus parity error(Command phase)
 - The file receives INITIATOR DETECTED ERROR MESSAGE during or at the conclusion of an information transfer phase (Command, Data In, Data Out or Status Phase)

Note: The initiator may send the INITIATOR DETECTED ERROR message as a result of an initiator detected SCSI Bus parity error or an internal error.
- Retry MESSAGE IN phase
 - The retry will be caused by the receipt of a MESSAGE PARITY ERROR message immediately following a MESSAGE IN phase.

Note: The Initiator may send the MESSAGE PARITY ERROR message as a result of an Initiator detected SCSI Bus parity error during the Message In phase.
- Receipt of multiple Identify message
 - The initiator is allowed to send a multiple IDENTIFY message out in order to toggle the disconnect/reconnect permission bit. This may be used to selectively enable or disable disconnect/reconnect permission during portion of a command. Note that this function does not effect the operation of the Forced Disconnect function.
- MESSAGE REJECT during Target Disconnection
 - If the Initiator rejects the SAVE DATA POINTER message, the file will disable disconnect/reconnect permission. This is equivalent to receiving an IDENTIFY message with bit 6 equal to zero. This will cause to file to inhibit the pending disconnection.
 - If the initiator rejects the DISCONNECT message, the file will not disconnect but may attempt to disconnect at a later time. This function may be used to selectively disable disconnection during portions of a command.

9.3 Attention Condition

The attention condition allows an initiator to inform the file that a MESSAGE OUT phase is desired. The initiator may create the attention condition by asserting the ATN signal at any time except during the ARBITRATION or BUS FREE phases.

The initiator must create the attention condition by asserting the ATN signal least two deskew delays before releasing ACK for the last byte transferred in a bus phase to guarantee that the attention condition will be honoured before transition to a new bus phase. This will guarantee a predictable file response to message received during the MESSAGE OUT phase for this attention condition. If the ATN signal is asserted later, it might be honoured in the current bus phase or the next bus phase and then may not result in the expected action.

After the initiator asserts the ATN signal, the file will respond with the MESSAGE OUT phase as follows:

Current Phase	Response
COMMAND	Message Out phase will occur after part or all of the Command Descriptor Block has been transferred to the file. The initiator must continue REQ/ACK handshakes during the Command phase until the file enters the MESSAGE OUT phase.
DATA	<p>The MESSAGE OUT phase will occur after part or all of the data bytes have been transferred and not necessarily on a logical block boundary. The initiator must continue REQ/ACK handshakes (asynchronous transfer) until it detects the phase change.</p> <p>Note: In synchronous transfer, the initiator must continue sending ACK pulses to reach an offset of zero.</p>
STATUS	The MESSAGE OUT phase will occur after the REQ/ACK handshake of the status byte has been completed.
MESSAGE IN	The MESSAGE OUT phase will occur before the file sends another message.
SELECTION	If ATN occurs during a SELECTION phase and before the initiator releases the BSY signal, the MESSAGE OUT phase will occur immediately after that SELECTION phase.
RESELECTION	The MESSAGE OUT phase will occur after the file has sent its IDENTIFY message for that RESELECTION phase. (First the file tries to complete the reselection.)

The initiator must keep the ATN signal asserted if more than one message byte is to be transferred during the MESSAGE Out phase. The file will process each message byte (multiple-bytes for an extended message) prior to receive the next message from the initiator. The file will continue to handshake and process byte(s) in the MESSAGE OUT phase until ATN goes false unless one of the following condition occurs:

1. The file receives an illegal or inappropriate message and goes to the MESSAGE IN phase to send a MESSAGE REJECT message.
2. The file detects a catastrophic error condition and goes to the BUS FREE phase.

9.4 SCSI Bus Related Error Handling Protocol

This protocol is used to handle error that threaten the integrity of a connection between the Target and an Initiator.

9.4.1 Unexpected BUS FREE Phase Error Condition

There are several error conditions that will cause the file to immediately change to the BUS FREE phase, regardless of the state of the ATN signal. The file will not attempt to reconnect to the initiator to complete the operation that was in progress when the error condition was detected. The initiator should interpret this as a catastrophic error condition.

If the LUN was identified by the file prior to the error condition, then the file will abort the active command for this initiator/LUN and generate sense data for this initiator/LUN to describe the cause of the catastrophic error. The initiator may retrieve this sense data by issuing a REQUEST SENSE command to this LUN. Note however, that the REQUEST SENSE command may fail if the catastrophic error condition persists.

If the LUN was not identified by the file prior to the error condition, then the file will not affect the sense data or the operation of any currently executing command for this initiator or any other initiator.

9.4.2 MESSAGE OUT Phase Parity Error

If the file detects a parity error during the MESSAGE OUT phase, the file will retry the MESSAGE OUT phase one time as follows:

1. Continue the REQ/ACK handshakes until the initiator drops ATN. The file will ignore all the remaining MESSAGE OUT phase bytes received after the parity error.
2. Assert the REQ signal prior to changing to any other phase. After detecting this condition, the initiator must resend all of the previous message byte(s) sent during this MESSAGE OUT phase.
3. Repeat the transfer of the MESSAGE OUT phase bytes(s). If the file receives all of the message byte(s) successfully, the file will change to:
 - Any other information transfer phase and transfer at least one byte.
 - BUS FREE phase if the message received was ABORT or BUS DEVICE RESET.

If a second parity error is detected, the target will abort the current command with CHECK CONDITION status and sense data of ABORTED COMMAND / SCSI PARITY ERROR .

9.4.3 MESSAGE IN Phase Parity Error (Message Parity Error)

If the file receives a MESSAGE PARITY ERROR message, it is considered a retrievable error. The file will do the following one time if no previous retrievable error and if a MESSAGE IN phase has just occurred:

1. Change phase to MESSAGE IN.
2. Send the last message again.

If this is the second retrievable error, the file will terminate the current command as follows:

1. Change to the BUS FREE phase, regardless of the state of the ATN signal.
2. Abort the active command for this initiator/LUN and set the sense data to ABORTED COMMAND / SCSI PARITY ERROR.

9.4.4 COMMAND Phase Parity Error

1. Change phase to MESSAGE IN and send a RESTORE POINTERS message.
2. If RESTORE POINTERS message is accepted, then change phase to COMMAND OUT and receive the command again.

If a second parity error is detected, the file will abort the current command with CHECK CONDITION status and sense data of ABORTED COMMAND / SCSI PARITY ERROR .

9.4.5 DATA OUT Phase Parity Error

If the file detects a parity error during DATA OUT phase, the file will abort the current command with CHECK CONDITION status and sense data of ABORTED COMMAND / SCSI PARITY ERROR.

9.4.6 INITIATOR DETECTED ERROR Message

An INITIATOR DETECTED ERROR message is valid after a COMMAND, DATA IN/OUT or STATUS phase has occurred. If any other phase has occurred, the message is rejected.

The recovery for COMMAND and STATUS phase consists of:

1. Change phase to MESSAGE IN and send a RESTORE POINTERS message
2. Repeat previous information phase .

If a second INITIATOR DETECTED ERROR message is received in the same selection, the target will abort the current command with CHECK CONDITION status and a Sense key of ABORTED COMMAND with additional sense code of INITIATOR DETECTED ERROR.

If a INITIATOR DETECTED ERROR message is received during a DATA IN/OUT phase, the target will abort the current command with CHECK CONDITION status and a Sense key of ABORTED COMMAND with additional sense code of INITIATOR DETECTED ERROR.

9.4.7 MESSAGE REJECT Message

The file will take the following actions after receiving the MESSAGE REJECT message in response to messages listed below.

DISCONNECT The file will not disconnect but remains connected.

COMMAND COMPLETE No error, continue to bus free.

IDENTIFY Command aborted - bus freed - Sense data set to MESSAGE REJECT ERROR.

LINKED CMD CMPLT Command aborted - link broken - bus freed - sense data set to MESSAGE REJECT ERROR.

MESSAGE REJECT Command aborted - STATUS phase executed with CHECK CONDITION - sense data set to MESSAGE REJECT ERROR.

RESTORE POINTERS Command aborted - status set to CHECK CONDITION - sense will be set with the error that caused the RESTORE POINTERS message to be issued. (Assuming that error recovery is in progress)

SAVE DATA POINTER The file will not disconnect from the SCSI bus. It will not be considered an error.

No previous Msg The command is aborted, the bus freed, and Sense data is set to MESSAGE REJECT ERROR. This occurs when the file has not sent a message, but gets a MESSAGE REJECT from the initiator.

10.0 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.

10.1 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.

10.1.1 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, 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 10.1.2, “Invalid LUN in Identify Message”)
2. Check Condition status for Incorrect Initiator Connection(see 10.1.3, “Incorrect Initiator Connection” on page 144)
3. Busy Status or Queue Full Status, or add command to command queue(see 10.1.4, “Command Processing During Execution of Active I/O process” on page 144)
4. Check Condition status for Unit Attention condition(see 10.1.5, “Unit Attention Condition” on page 147)
5. Check Condition status during Start-up and Format operations(see 10.1.6, “Command Processing During Start-up and Format Operations” on page 147)
6. Check Condition status for Internal Error Condition(see 10.1.7, “Internal Error Condition” on page 148)
7. Check Condition status for Deferred Error Condition(see 10.1.8, “Deferred error” on page 149)
8. Check Condition status for Degraded Mode Condition(see 10.1.9, “Degraded Mode” on page 149)
9. Reservation Conflict status(see 10.1.11, “Command Processing While Reserved” on page 150)
10. Check Condition status for invalid command opcode
11. 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.

10.1.2 Invalid LUN in Identify Message

There are three different circumstances defined within the SCSI protocol when the response to an invalid LUN will occur. Each of these result in a different response.

10.1.2.1 Case 1 - Selection message sequence with Inquiry command

The INQUIRY command is a special case in SCSI. It is used to configure the bus when file IDs and LUNs are not known. The proper response is to return the inquiry data with a peripheral drive type of 1Fh which indicates that the specified LUN is not supported.

10.1.2.2 Case 2 - Selection message sequence with any other command

Any other commands, except REQUEST SENSE, return CHECK CONDITION status when an invalid LUN is specified in the message sequence following selection. In response to a REQUEST SENSE command the target shall return sense data. The sense key shall be set to ILLEGAL REQUEST and the additional sense code shall be set to LOGICAL UNIT NOT SUPPORTED.

10.1.2.3 Case 3 - After selection message sequence

It is permissible for the initiator to issue multiple IDENTIFY messages during a single command sequence provided to LUN remains the same. If the LUN is altered, the file goes to a Bus Free Phase.

10.1.3 Incorrect Initiator Connection

It is an Incorrect Initiator Connection error if any of the following occurs:

- an Initiator attempts to establish an I_T_L nexus when an I/O process (either queued or active) with an I_T_L nexus already exists from a previous connection with the same initiator.
- an Initiator attempts to establish an I_T_L_Q nexus when an I_T_L nexus already exists from a previous connection with the same initiator.
- an Initiator attempts to establish an I_T_L nexus when an I_T_L_Q nexus already exists from a previous connection with the same initiator.

Note: It is not an Incorrect Initiator Connection to send a Request Sense command without a Queue tag message when sense is pending on the logical unit for the Initiator that issues the Request Sense command.

- an Initiator attempts to establish an I_T_L_Q nexus when an I/O process (either queued or active) with an I_T_L_Q nexus already exists from a previous connection with the same initiator.

If any of the above errors occur, all queued I/O processes and active I/O processes associated with the issuing Initiator on the specified logical unit are terminated. The current I/O process is ended with a CHECK CONDITION status, the sense key is set to ABORTED COMMAND and the additional sense code is set to OVERLAPPED COMMANDS ATTEMPTED. Status is only returned for the current I/O process.

10.1.4 Command Processing During Execution of Active I/O process

When the Target is not executing any active I/O processes, a new I/O process is permitted to execute (unless execution is prevented by another internal Target condition listed in 10.1.1, “Priority of SCSI Status Byte Reporting” on page 143).

If an active I/O process does exist when the Target receives a new command, then the Target determines if:

- Check Condition Status with Sense Key = Aborted Command is returned for an Overlapped Commands Attempted error
- the command is permitted to execute
- the command is added to the command queue
- Queue Full Status is returned
- Busy Status is returned

If an active I/O process does exist 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 Aborted Command for an Overlapped Commands Attempted error if:
 - See 10.1.3, “Incorrect Initiator Connection” on page 144
- the command is permitted to execute if
 - the command is an Inquiry or Request Sense command
- Check Condition Status is returned with Sense Key set to Logical Unit Not Ready if:
 - the start-up operation or format operation is an active process.
- the command is permitted to execute if
 - the conditions to execute concurrently are met. (See 10.5, “Concurrent I/O Process” on page 156)
- the command is added to the command queue for an I_T_L nexus if:
 - no Queue Tag message was received during the connection which established the I/O process, and
 - Untagged Queuing is enabled(UQE = 1), and
 - disconnection is allowed for the current I/O process, and
 - there is no queued I/O process or active I/O process corresponding to the I_T_L nexus for the current I/O process, and
 - the command is not linked to a previous command.
- the command is added to the command queue for an I_T_L_Q nexus if:
 - a Queue Tag message was received during the connection which established the I/O process, and
 - Tagged Queuing is enabled(DQue = 0), and
 - an I/O process(either active or queued) exists at the Target for this Initiator, and
 - disconnection is allowed for the current I/O process, and
 - there is no queued I/O process or active I/O process corresponding to the I_T_L_Q nexus for the current I/O process, and
 - the command is not linked to a previous command.

Note: Both Tagged and Untagged Queuing must be enabled(DQue = 0 and UQE = 1) for the Target to the queue tagged I/O processes from multiple Initiators.

- Queue Full Status is returned if:
 - the command would otherwise be queued(according to the rules described above)but the command queue is full and all slots are utilized, or
 - the command would otherwise be queued(according to the rules described above)but all of the available command queue slots not reserved for use by another initiator are utilized, or
 - Tagged Queuing is enabled(DQue = 0) and a Format Unit command was previously queued but has not yet begun execution, or
 - Tagged Queuing is enabled(DQue = 0) and a Start Unit command was previously queued but has not yet begun execution.
- Busy Status is returned if:
 - Tagged Queuing is disabled(DQue = 1) and a Format Unit command was previously queued but has not yet begun execution, or
 - Tagged Queuing is disabled(DQue = 1) and a Start Unit command was previously queued but has not yet begun execution, or
 - the command would otherwise be queued(according to the rules described above)but disconnection is not allowed for the current I/O process, or
 - the command would otherwise be queued(according to the rules described above)but Untagged Queuing is disabled(UQE = 0) and an I/O process (either active or queued) exists at the Target from a different Initiator.

If a command is queued, command execution may still be prevented at a later time when the command is dequeued to become an active I/O process. This occurs if command execution is prevented by another internal Target condition listed in 10.1.1, “Priority of SCSI Status Byte Reporting” on page 143 at the time the command is dequeued.

10.1.5 Unit Attention Condition

The file will generate a unit attention condition for each initiator whenever:

- The file has been reset.
This includes Power On Reset, SCSI Bus Reset, SCSI BUS DEVICE RESET message.
- The mode parameters in effect for this initiator has been changed by another initiator.
- The microcode has been changed.
WRITE BUFFER command has been executed to download microcode. In this case, a unit attention condition is generated for all initiators except the one that issued the command.
- Commands are cleared by another initiator.
This condition is generated against the initiator that has queued commands, if ...
 - Clear Queue Message is received.
 - Contingent Allegiance Condition is cleared when QERR (in Mode Page 0A) is 1.
 - DQue is set to 1 while queued command exist.

The unit attention condition persists for each initiator until that initiator clears the condition as described in the following paragraphs.

If the file receives a command from each initiator before reporting a CHECK CONDITION status for a pending unit attention condition for that initiator , the file's response varies with the command as follows.

INQUIRY The file executes the command with GOOD status and preserves the unit attention condition.

REQUEST SENSE

If the file has an available pending sense data for the initiator, the file sends the pending sense data and preserves the unit attention condition for the initiator.

If the file does not have an available pending sense data for the initiator, the file sends sense data for the unit attention condition and clear the unit attention condition for the initiator.

ALL OTHER The file terminates the command with a CHECK CONDITION status and preserve the unit attention condition.

If the file receives a command form each initiator after reporting a CHECK CONDITION status for a pending unit attention condition for that initiator , the file's response varies with the command as follows.

REQUEST SENSE The file sends the sense data for a pending unit attention condition and returns GOOD status. And the file clear the unit attention condition for the initiator.

ALL OTHER The file execute the command with GOOD status and clear the unit attention condition unless another unit attention condition exists. And then the sense data for the unit attention condition is lost.

10.1.6 Command Processing During Start-up and Format Operations

If the Target receives a command from an Initiator while the Target is executing a start-up or format operation, The Target's response varies with the command as follows:

INQUIRY The file sends a inquiry data and returns appropriate status.

REQUEST SENSE Execute the command, returns a Sense key of NOT READY and an Additional Sense Code of LOGICAL UNIT NOT READY and return GOOD STATUS.

The Additional Sense Code Qualifier that is returned dependent on type of I/O processes that are active:

For the START/STOP UNIT 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.

START/STOP UNIT If Untagged Queuing is enabled(UQE = 1), and the start-up operation is an active process and a Start/Stop Unit command(either active or queued) does not exist at the Target from this initiator, and disconnection is allowed for the current I/O process then: The command is added to the command queue.

Otherwise: Do not execute the command and return Check Condition Status. The Sense data generated is described in Request Sense above.

ALL OTHER The file terminates the command with CHECK CONDITION status. The Sense data generated is described in Request Sense above.

10.1.7 Internal Error Condition

The Target generates an Internal Error condition for all Initiators when:

- an internally initiated operations with an unrecoverable error. the following is a list of internally initiated error conditions:
 - During the execution of the start-up sequence for Auto Start after the SCSI bus has been enabled and prior to completion of the bring-up sequence.
 - Following a SCSI H/W reset or a SCSI Bus Device Reset message if the reset was received during a start-up sequence with the Auto Start function enabled. The start-up sequence is executed if it has not been previously executed and completed.
- an recoverable error occurs during an internal Target idle time function

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. Any outstanding Deferred Error condition is cleared for all initiators and the associated Sense data is lost.

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 Target's response varies with the command as follows:

INQUIRY The file executes the command with GOOD status and do not clear the Internal Error condition.

REQUEST SENSE The file executes the command, return the sense data generated by the Internal Error condition, return Good Status, and clear the Internal Error condition for that Initiator.

ALL OTHER The file terminates the command with a CHECK CONDITION status and clear the Internal Error condition.

10.1.8 Deferred error

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 Start/Stop Unit command with the immediate bit of one ends with an error.
- Execution of a Format Unit command with the immediate bit of one ends with an error.
- Execution of a Write command with WCE (write cache enable) bit of one ends with an error.

10.1.9 Degraded Mode

There are certain errors or conditions which may impair the file's ability to function normally. Rather than fail hard, the file 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.

10.1.10 Degraded mode handling

If any of degrade mode condition occurs, an initiator can receive the sense key of Hardware error (4h). After clearing the sense key, the drive accepts commands according to the status of the degrade mode.

The following table shows the degrade mode status, acceptable command and additional sense code. If the degrade mode exists, the Hard ware error condition caused by the degrade mode will be recreated any time after POR, SCSI Reset or Bus Device Reset message. Following list shows the various operation mode.

10.1.10.1 Degrade Mode Entry Condition/Consequence State

Degrade-Mode	Entry Reason	Accepting Request	Description and sense code
Power on Self Test fail	<ul style="list-style-type: none"> • Failure of a Send Diagnostic self-test • Failure of a start-up sequence 	<ul style="list-style-type: none"> • Request Sense • Inquiry • Start/Stop unit • Write Buffer(except download and save) 	<p>The integrity of the drive is questionable.</p> <p>Sense code = 4080 (diagnostic fail), Sense code = 4081 (HDC), 4082 (HIC), 4083 (other) 4084 (RAM error)</p>

Degrade-Mode	Entry Reason	Accepting Request	Description and sense code
Spin-up	Can not start up motor	<ul style="list-style-type: none"> Request Sense Inquiry Start/Stop unit Write Buffer(except download and save) 	Return Spin-up Degrade Mode. Sense code = 0400
U-code	Can not read u-code	Same as Spin-up Degrade	Return U-code Degrade Mode. Sense code = 4085 (u-code download fail)
Reserved Area	<ul style="list-style-type: none"> Fail to read Push Table Fail to read saved Mode Parameter Fail to read Defect list 	All commands.	Reserved area sector valid check failed. Defective sector found in reserved area. (Hard Error) The integrity of the drive is questionable. Sense code = 4080 (diagnostic fail), 1900 (Defect list error),
Config	Fail to read Drive Configuration Table	All commands	Configuration sector valid check failed. The integrity of the drive is questionable. May return blanks in Inquiry Command. Sense code = 4C00 (Self configuration fail)

10.1.11 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:

1. If the issuing Initiator is the one that made the reservation and also the one to receive the reservation then:
 - All commands are permitted.

2. 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.
3. 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.
4. If the issuing Initiator is not the one that made the reservation but is the one to receive the reservation then:
 - An 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 10.1.1, “Priority of SCSI Status Byte Reporting” on page 143

10.2 Priority Commands

Certain SCSI commands always execute without returning a Busy Status, Reservation Conflict Status in response to the command. These commands are:

- Inquiry
- Request Sense

These commands do not disconnect from the SCSI bus prior to completion. They are executed prior to attempting to complete the execution of any other pending command that has disconnected from the SCSI bus. Therefore, a second priority command cannot be received during the execution of a priority command.

These commands are never queued whether or not the command is sent with a queue tag. However, the rule for an Incorrect Initiator Connection still apply to priority commands.(see 10.1.3, “Incorrect Initiator Connection” on page 144)

10.3 Command queuing

When the initiator specifies that the file shall disable command queuing, the initiator must send only untagged command. 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 and will be immediately executed without Bus disconnection

- Priority Commands(i.e.: Request Sense and Inquiry)
- Commands linked to previous commands. These are defined to be part of a single I/O process. (Linked commands are always executed immediately following the previous command from the same initiator. No other

Initiator's command are allowed to be executed between two linked commands.)

- Commands for which disconnection is not allowed.(These may result in a Busy Status.)
- Commands in which a SCSI bus error occurred between selection and first disconnection following the receipt of the CDB.
- Commands for an invalid LUN.
- Commands which cause an OVERLAPPED COMMANDS ATTEMPTED error(see 10.1.3, “Incorrect Initiator Connection” on page 144)

10.3.1 Queue depth

The file has 32 queue elements. Seven of these are reserved for the seven possible initiator's that may connect to the file. This leaves a pool of 25 queue elements which are shared by all initiators on a first come first served basis. 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 26 commands, if no other initiator has more than one command in the queue, although at times this maximum may be reduced as the file can reserve command blocks for internal use.

10.3.2 Tagged queuing

Commands with a tag message are saved in the command queue. Queued commands will be reordered in the method of 'Single side elevator'. See the section on Reordering 10.4, “Command reordering” on page 154 for details.

10.3.3 Untagged queuing

When Untagged Queuing is enabled, the target supports queuing one I/O process from each initiator. If the target receives an untagged I/O process while executing an I/O process from a different initiator, the untagged I/O process may be queued.

Untagged I/O process are treated by the target as though they were received with Simple Queue Tag messages for purposes of queuing.

Note: There is no guarantee that I/O processes are executed in the order they were received in a multiple initiator environment when Untagged Queuing is enabled.

10.3.4 Command queuing rule

Commands can be received during an active I/O process if Bus is free. If CPE (concurrent process enable) bit of Mode page 0 is 0, only a single command except 2 commands (Inquiry and Request Sense) can be executed at the same time. In that case, a command starts to be executed after the previous command has completed. If CPE (concurrent process enable) bit of Mode page 0 is 1, I/O process of Read(6), Read extend(10), Write(6) and Write extend(10) can be active at the same time. See 10.5, “Concurrent I/O Process” on page 156 for details.

When the target receives a new command, the target do the following things.

1. Check if the command is overlapped. A target that detects an Incorrect Initiator Connection shall abort all (queued and active) I/O processes for the initiator and shall return Check Condition status. The sense key shall be set to Abort Command (0Bh) and the additional sense code shall be set to Overlapped Commands Attempted (4E00h).

An incorrect initiator connection occurs when an initiator:

- a. attempts to establish an I_T_L nexus when an I_T_L nexus already exists from a previous connection.
(For example, an initiator sends an untagged command when an untagged command is queued or being currently executed for the initiator.) or,
- b. attempts to establish an I_T_L_Q nexus when an I_T_L nexus already exists from a previous connection.
(For example, an initiator sends a tagged command when an untagged command is queued or being currently executed for the initiator.) or,
- c. attempts to establish an I_T_L nexus when an I_T_L_Q nexus already exists from a previous connection.
(For example, an initiator sends an untagged command when a tagged command is queued or being currently executed for the initiator.) or,
- d. attempts to establish an I_T_L_Q nexus when the same I_T_L_Q nexus already exists from a previous connection.
(For example, an initiator sends a tagged command when a tagged command with the same Queue Tag is queued or being currently executed for the initiator.)

It is not an Incorrect Initiator Connection to reconnect to an already established I_T_L or I_T_L_Q nexus if the initiator sends an Abort, Abort Tag, Clear Queue, Bus Device Reset during the same Message Out phase as the Identify message.

2. Check if the command is valid.
 - a. Check if the command code is valid. If the command code is invalid, the target sends Check Condition and sets the sense key Illegal request (05h) and the additional sense code Invalid command (2000h).
 - b. Check if LINK = 0 and FLAG = 1. If LINK = 0 and FLAG = 1 then the target sends Check Condition and sets the sense key Illegal request (05h) and the additional sense code Invalid CDB (2400h).
 - c. Check if parity error occurred on CDB receipt. If parity error occurred, the target sends Restore pointer message (03h) to retry. If parity error occurred again, the target sends Check Condition and sets the sense key Aborted Command (0Bh) and the additional sense code SCSI Parity (4700h).
 - d. Check if the reserved bits are off. If any reserved bit in CDB is on, the target sends Check Condition and sets the sense key Illegal Request (05h) and the additional sense code Invalid CDB (2400h).
3. Check if the command must be immediately executed. If the command is one of either Inquiry or Request Sense sent as untagged and unlinked command, it must be executed immediately without disconnecting Bus.
4. Check if the queue is full. If the command queue is full, the target sends Queue Full Status (28h).
5. Check if Busy. Busy Status (08h) is returned if:
 - disconnection is not allowed for the incoming tagged command.
 - disconnection is not allowed when queued or active process exists.
6. If all above is satisfied, add the command into the queue.

10.3.5 Queue Full status

This status is returned when a Simple Queue tag, Ordered Queue tag or Head of Queue tag message is received and the command queue is full. The I/O process is not placed in the command queue. Since one queue element is reserved for each initiator, any untagged command that doesn't cause Incorrect Initiator Connection will not cause Queue Full status.

10.3.6 Device behaviour on Command queuing

1. Initiators must send a Queue tag immediately after the Identify message in Message Out phase just after Selection. Targets send a simple queue tag immediately after the Identify message in Message In phase just after Reselection.
2. Each initiator can issue either one of tagged command or untagged command exclusively at the same time. There can exist other initiator using the other.
3. When DQue (Disable queue) of mode page 0Ah is 1, if an initiator issues a tagged command, the drive returns "Message Reject" message (07h) and receives that command as an untagged command.
4. Queue Tag number doesn't affect the order to execute.
5. An initiator issues a command with a queue tag which is same as the current I/O process or queued I/O process, the target returns Incorrect Initiator connection.
6. A series of linked commands are a single I/O process, and are assigned the queue tag established in the initial selection. A command received with a Head-of-Queue tag message shall not suspend a series of linked commands for which the target has begun execution.
7. If DQue is changed to 1 while queued commands exist, all queued commands for the ITL nexus issuing the mode select will be aborted. All future commands received, from any initiator, with a queue tag will be processed as untagged commands, with a message reject message being returned immediately after the qtag is received by the target.

10.4 Command reordering

Command reordering function is supported under tagged command queuing enabled (DQue = 0). Reorder feature reorders Read/Write commands in order to minimize seek time between commands. This function will improve total thruput of the drive.

10.4.1 Reorder method

The reordering algorithm of the drive is 'Single side elevator'. The queued commands are reordered as ascending order of the requesting LBA from the LBA of the current command. Commands of smaller LBA will be queued after the commands of greater LBA.

For example, suppose that a Read command requesting LBA 1000 is being executed. And the commands queued are ...

- Read of LBA 100.
- Read of LBA 400.
- Read of LBA 1200.
- Read of LBA 1400.

These commands must be reordered ...

1. Read of LBA 1200.
2. Read of LBA 1400.
3. Read of LBA 100.
4. Read of LBA 400.

While executing the current command, if the incoming commands are Read of LBA 1500, Read of LBA 120, and Read of LBA 35, after the commands are received, the new queue becomes ...

1. Read of LBA 1200.
2. Read of LBA 1400.

3. Read of LBA 1500.
4. Read of LBA 35.
5. Read of LBA 100.
6. Read of LBA 120.
7. Read of LBA 400.

Reordering is done each time a command is received. The order of the commands previously received, reordered and queued will not be changed. Reordering is done only by inserting a incoming command into the best place in the command queue.

10.4.2 Command aging

If the command aging function of the drive is enabled, via mode page 0h, then the re-ordering function of the queue will be modified based on the length of time a command has remained in the queue.

Until the user specified time limit a command will remain in the queue and may or may not be executed based on the other commands in the queue at the time and their relative locations - as described above.

Once the time limit for a command has expired it will be executed next, i.e. when the drive has completed the current command in progress. If more than one command expires then these commands will be executed in the order they were received by the target.

10.4.3 The restriction of reordering

10.4.3.1 SCSI restriction

10.4.3.1.1 Tag message: The initiator can send an untagged command or 3 kinds of tagged message to define reordering manner.

- Untagged command.
The initiator does not care the reordering feature. An untagged command might be passed by a command sent by another initiator. Untagged command will not create data integrity problem by the same initiator, because the initiator cannot send more than one untagged command at the same time.
- Simple Queue Tag message (20h).
This tagged command can be reordered by the drive unless it violates the rule of the Head-of-queue tag and the Ordered-queue tag message.
- Head of Queue Tag message (21h).
Commands with this tag should be inserted into the head of the tag queue. While a command is being executed, this tagged command will be inserted to the head of queue to be executed after the command being currently executed. The previous command being executed will not be terminated by this tagged command. If the plural head-of-queue tagged commands are received, those commands will be executed in LIFO (Last in First out) order.
- Ordered Queue Tag message (22h).
This tagged command is executed in the order received. All commands received before this command should be executed before this command, and all commands received after this commands should be executed after this command.

10.4.3.1.2 Mode parameter (Control Mode page 0Ah): Queue Algorithm modifier (7-4 bit of byte 3) specifies restriction on the algorithm used for reordering commands that are tagged with Simple-queue tag message.

0h : Restricted reordering. The target shall order the actual execution sequence of the queued commands from each initiator such that data integrity is maintained for that initiator.

- 1h : Un-restricted 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 the command sequence order are explicitly handled by the initiator.
- 2h-7h : Reserved. The drive accepts these values, but the operation is not defined.
- 08h Command reordering is not performed.
- 9h-Fh : Reserved. The drive accepts these values, but the operation is not defined.

10.4.3.2 Implementation restriction

There are some restrictions that are not clearly stated in SCSI specifications, but they are implemented in the drive.

1. Priority commands. Untagged and unlinked Request sense and Inquiry
These are executed immediately after receiving CDB without disconnecting SCSI bus. They do not follow the rule of reordering. The execution of these command will not affect data integrity.
2. Commands other than Read/Write.
Commands such as Mode Select, Reserve, Release or Start Stop Unit may affect the condition of execution or data of commands that follows. If Queue algorithm modifier is 0 (Restricted reordering), all commands including these commands except Read/Write and priority commands (untagged and unlinked Request sense and Inquiry) will not be reordered even if untagged or Simple Queue tagged. In that case, these commands will be always treated as Ordered tagged command.
3. Concurrent I/O.
Commands with Head of Queue tag or Ordered Queue tag will not be executed concurrently. Only untagged and Simple-queue tagged Read/Write command can be executed concurrently.

10.4.4 Reordering general rules

1. Reordering follows SCSI restrictions and implementation restrictions.
2. If queue algorithm modifier = 1 (Unrestricted reordering), untagged and Simple-queue tagged Read and Write commands will be reordered in 'Single side elevator' method of LBA, regardless of breaking data integrity. Non Read/Write commands will be passed by a Read/Write command unless it violates tag message restrictions.
3. If queue algorithm modifier = 0 (Restricted reordering), untagged and Simple-queue tagged Read and Write commands will be reordered in 'Single side elevator' method of LBA, as far as it maintains data integrity. It implies that a Read command whose requested LBAs overlap those of a Write command previously received will not pass the Write command. A Write command will be the same. Non Read/Write commands will not be passed by any Read/Write or non Read/Write command.

10.5 Concurrent I/O Process

The Concurrent I/O process is that plural I/O processes are active (not queued) on the same logical unit at the same time. The target may starts the data phase of an I/O process while another I/O process is not completed. The following I/O processes are allowed to executed concurrently.

- Unlinked and untagged Request Sense and Inquiry during execution of other commands.
- When CPE (Concurrent process enable) bit is 1, one of the following commands can be executed during another one or the same one of the following commands is being executed, if those are untagged or simple tagged commands.
 - Read(6), Read extend(10)
 - Write(6), Write extend(10)

When an I/O process ends in Check Condition Status, the drive enters the Contingent Allegiance Condition and other queued I/O processes from all initiators on the same logical unit will not reconnect and will not complete the execution until the sense data is cleared. See 10.17, “Contingent allegiance Condition” on page 167 for details. If an I/O process (P-1) encounters an error while another I/O process (P-2) is active, the drive returns Check Condition to P-1 and P-2. The drive may continue P-2 until its convenient point to suspend, but may not send a Status. After the initiator clears the Contingent Allegiance condition, the drive will resume or terminate P-2 according to QErr bit of Control mode page.

If the drive gets an error of P-2 before suspending the execution of P-2, it will keep the sense data separately from the sense data for P-1. The sense data for P-2 will be set after the Contingent Allegiance condition caused by P-1 is cleared.

10.6 Back to Back Write

Back to Back Write allows plural write commands requesting sequential LBA to write without losing a motor revolution.

The target will write back to back under the following conditions:

- Two Write(6) or Write extend(10) commands executed consecutively.
- The write commands address consecutive LBA. The LBA specified by the 2nd write command is the next to the last LBA by the 1st write command.
- The data of the LBA for the 2nd write command is received before the last LBA of the 1st write command is written to the media.
- The CPE (Concurrent process enable) bit is 1 or the WCE (Write Cache enable) bit is 1.

If the CPE (Concurrent process enable) is 1 and the WCE (Write cache enable) is 0, the Data phase of the 2nd write command is started after the Data phase of the 1st write command but before the Status phase of it. After seek complete, the data of both 1st and 2nd write command will be written onto the media. Then the drive reconnects to the nexus of the 1st write command to send Status before sending Status to the 2nd write command.

If the CPE (Concurrent process enable) is 0 and the WCE (Write cache enable) is 1, the drive returns Good Status immediately after the Data phase of the 1st write command but before actually writing the data to the media. The I/O process of the 2nd write command begins after the drive sends Good Status to the 1st write command. And then Good Status is returned to the 2nd write command. After seek complete, the data of both 1st and 2nd write command will be written onto the media.

If the drive returns Check Condition to the 1st write command and enters the contingent allegiance condition, execution of the 2nd write command is suspended. The 2nd write command will resumed execution when the contingent allegiance condition is cleared.

10.7 Write Cache

If the WCE (Write cache enable) bit is 1, the drive returns Good Status and Command Complete message and goes to Bus Free immediately after receiving the data of the last sector before actually writing the data onto the media. The drive will accept and queue a command, but it can not start to execute the command after sending a Good Status except the following condition.

- The incoming command is one of priority commands.
- The CPE (Concurrent process enable) is 1 and the incoming command is Read(6), Read extend(10), Write(6) Write extend(10).

The drive behavior according to the mode parameter.

When Write Cache is enabled and Concurrent I/O process is enabled,
the drive may start and complete following Read/Write commands before the actual write operation (writing on the media) is completed. So, any following Read/Write commands or priority commands returning a Good Status. do not guarantee the completion of the Write command. Under the current implementation, a command except Read/Write commands or priority commands returning a Good Status following a Write command can guarantee that the data is written to the media.
A Synchronize Cache command always performs this function regardless of the current implementation.

When Write Cache is enabled and Concurrent I/O process is disabled,
the drive may not start any following commands except priority commands before a previous Write command is completed and the drive sends a Status. So, under the current implementation, any following commands except priority commands returning a Good Status guarantees that the data is written to the media.
A Synchronize Cache command always performs this function regardless of the current implementation.

When Write Cache is disabled and Concurrent I/O process is enabled,
the drive may start and complete following Read/Write commands before a previous Write command is completed. But the Write command will not be completed before the data is written to the media.
A Synchronize Cache command is not needed in order to assume that the data is written to the media.

When Write Cache is disabled and Concurrent I/O process is disabled,
the drive may not start any following commands except priority commands before a previous Write command is completed and the drive sends a Status. And the Write command will not be completed before the data is written to the media.
A Synchronize Cache command is not needed in order to assume that the data is written to the media.


If the drive detects an error after it returns a Good Status, the drive sets a Differed 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.

10.8 Power Saving Mode

10.8.1 Overview

Power save function will save power consumption while the drive is idle. The drive automatically transfers its operating mode according to the event and timer.

10.8.1.1 Mode Transition



Mode Characteristics

Mode Name	Idle	Standby	Stopped
HDC	On	Off	Off
SRV	On	Off	Off
R/W Channel	Off	Off	Off
Content of sector buffer (Write Cache, Look Ahead & etc.)	Kept	Lost	Lost
Spindle motor	On	Off	Off
Status	Good	Good	Not Ready

Trigger of Mode Transition

To From	Idle	Standby	Stopped
Idle	---	Expiration of Standby Timer & Read/Write Idle	SCSI Stop Unit Command
Standby	Any SCSI Command	---	SCSI Stop Unit Command
Stopped	SCSI Start Unit Command	N/A	---

Action for Mode Transition

To From	Idle	Standby	Stopped
Idle	---	Act_Lock Channel_off VCM_off Stop_spindle	Act_Lock Channel_off VCM_off Stop_spindle
Standby	Channel_on VCM_on HDC_on Start_spindle	---	Not Ready Status
Stopped	Channel_on VCM_on HDC_on Start_spindle	N/A	---

10.9 Automatic Rewrite/Reallocate

The target supports Auto and Recommended Reallocate for READ, WRITE, WRITE and VERIFY, VERIFY.

Automatic and Recommend Reallocate operate from within the read/write command. When an automatic reallocation occurs, the read or write command takes longer to complete. This operation is sometimes called "Auto-Reassign" due to its similarity to Reassign command. During this time, the target disconnects from the SCSI bus, if allowed, and reconnects before ending the command.

Following is a description of the target behavior for each setting of ARRE. ARRE setting effects all data errors.(No Sector Found, Data Sync Byte Errors and Data ECC Errors.)

ARRE=1 : An error site determined to need rewriting or reassignment during a read is automatically rewritten or reallocated at the conclusion of the read and prior to sending the status. If the site cannot be automatically rewritten or reallocated, then a recommendation for reassign-

ment is given. The site will be automatically rewritten or reallocated only if the data has been successfully read.

ARRE=0 : An error site determined to need rewriting or reassignment during a read is recommended for rewriting or reassignment at the conclusion of the read.

The setting of the ARRE bit is checked and the target will automatically rewrite/reallocate or recommend rewrite/reassign for the following commands.

- Read(6)
- Read(10)

Target will recommend rewrite/reallocate but will not auto rewrite/ reallocate for the following commands.

- Verify
- Verify Portion of Write and Verify

For all other commands the ARRE setting is ignored and the target will not automatically rewrite/reallocate or recommend rewrite/reassign.

Following is a description of the target behavior for each setting of AWRE. AWRE setting effects only No Sector Found Errors on writes.

AWRE=1 : An error site determined to need reassignment during a write is automatically reallocated at the conclusion of the write and prior to sending the status. If the site cannot be automatically reallocated, then a recommendation for reassignment is given. The site will be automatically reallocated only if the write recovery succeeded at the conclusion of the write.

AWRE=0 : An error site determined to need reassignment during a write is recommended for reassignment at the conclusion of the write.

The setting of the AWRE bit is checked and the target will automatically reallocate or recommend reassign for the following commands.

- Write(6)
- Write(10)
- Write portion of Write and Verify

For all other commands the AWRE setting is ignored and the target will not automatically reallocate or recommend reassign.

Auto/Recommend Reallocate information is communicated via the sense data returned following a command during which a site was determined to need rewriting or reassignment. The LBA returned in the sense data is the LBA that determined to need rewriting or reassignment.

The sense data combinations with auto/recommend rewrite/reallocate are listed below.

Key	Code	Qual	Description
1	17	00	Recovered Data without ECC.
1	17	06	Recovered Data without ECC - Auto Reallocated.
1	17	07	Recovered Data without ECC - Recommend Reassign.

Key	Code	Qual	Description
1	17	09	Recovered Data without ECC - Data Rewritten.
1	18	00	Recovered Data with ECC.
1	18	02	Recovered Data with ECC - Auto Reallocated.
1	18	05	Recovered Data with ECC - Recommend Reassign.
1	18	07	Recovered Data with ECC - Data Rewritten.

10.10 Segmented Caching

10.10.1 Overview

Segmented Caching divides the data buffer into several smaller buffers. Each buffer is used as Read/Write/Read-Ahead buffer.

The size of segmented buffer is controlled by number of segmented buffer field of read cache page (page 08h). But file may adapt size of segmented buffer by checking access pattern.

10.10.2 Read Ahead

The Read Ahead function consists of reading data that the Initiator has not yet requested to the file buffer. This function is intended to improve performance for an initiator that frequently accesses sequential data with successive SCSI read commands. The Read Ahead function works when RCD (the read cache disable) bit of read cache page (page 08h) is set to 0.

The file will continue to read subsequent logical blocks following the rules below.

1. At least the same number of sectors succeeding the sectors requested by the last Read command will be read automatically if they are not available on the sector buffer.
2. The drive has at least 2 segmented buffers. Data stored by the most recent Read ahead function will not be discarded by Write command.
3. Data in the segmented buffers will not be replaced unless it is against the rule 1.
4. The Read Ahead operations continue across all physical boundaries (such as tracks and cylinders).

Commands that affects the Read ahead function.

The drive initiates the Read ahead function

- RCD is 0, and
- Read(6), Read extended(10), Read Verify, and Write and Verify is received, and
- The consecutive LBA of the requested LBA is not available on the buffer.

The action of each command for previously started Read ahead function is listed below. If SCSI reset or bus device reset message is received, all contents of segmented buffer is flushed.

Code	Command	Action
00h	Test Unit Ready	No Effect
01h	Rezero Unit	Abort Read Ahead
03h	Request Sense	No Effect
04h	Format Unit	Abort Read ahead and flush all segments
07h	Reassign Blocks	Abort Read ahead and flush all segments

Code	Command	Action
08h	Read(6)	Abort Read ahead if data not in active segment
0Ah	Write(6)	Abort Read ahead and flush the LRU segment
0Bh	Seek(6)	Abort Read Ahead
12h	Inquiry	No Effect
15h	Mode Select(6)	Abort Read ahead and flush all segments
16h	Reserve	No Effect
17h	Release	No Effect
1Ah	Mode Sense(6)	Abort Read ahead and flush all segments
1Bh	Start/Stop Unit	Abort Read ahead and flush all segments
1Dh	Send Diagnostic	Abort Read ahead and flush all segments
25h	Read Capacity	No Effect
28h	Read extended(10)	Abort Read ahead if data not in active segment
2Ah	Write extended(10)	Abort Read ahead and flush the LRU segment
2Bh	Seek extended(10)	Abort Read Ahead
2Eh	Write and Verify	Abort Read ahead and flush the LRU segment
2Fh	Verify	Abort Read ahead and flush the LRU segment
34h	Pre-Fetch	Abort Read ahead if data not in active segment
35h	Synchronize Cache	Abort Read ahead and flush all segments
37h	Read Defect Data	Abort Read ahead and flush all segments
3Bh	Write Buffer	Abort Read ahead and flush all segments
3Ch	Read Buffer	Abort Read Ahead

Code	Command	Action
3Eh	Read Long	Abort Read ahead and flush the LRU segment
3Fh	Write Long	Abort Read ahead and flush the LRU segment
4Dh	Log sense	Abort Read ahead and flush all segments

Figure 101. Read Ahead handling per each command

Even if an error occurs during Read ahead, the error will not be reported to the Initiator. The data read before the error occurs will be stored as a valid data by Read ahead.

10.11 Reselection

A reselection timeout error occurs when the target attempts to reselect an initiator and the initiator does not respond within a Selection Timeout delay(250 mSec). If this occurs, the target releases the SCSI bus to the Bus Free phase(following the reselection timeout procedure) and then retries the reselection up to 16 times. If a 16th reselection timeout error occurs, the target releases the SCSI bus to the Bus Free phase(following the reselection timeout procedure) and does not attempt to reselect the initiator again. This error condition causes the target to generate sense data with a Sense Key of Aborted Command and an Additional Sense code of Select/Reselect failure(45h). If a reselection is successful, command execution is resumed and no sense data is generated.

10.12 Single Initiator Selection

For single initiator systems, it is not an error to have only the target ID bit present during selection. Disconnection is not allowed for Single Initiator Selection with only one ID bit present during selection. The initiator must not send an Identify message with the disconnect permission bit(6) on.

10.13 Non-arbitrating systems

The Target cannot detect whether other SCSI devices on the SCSI bus use arbitration prior to selection. As a consequence, the Target allows disconnect permission to be enabled by the Identify message independent of the initiators use of arbitration prior to selection. A non-arbitrating initiator must ensure that disconnect permission in the Identify message is disabled (bit 6=0)for proper operation.

10.14 Selection without ATN

If the target is selected without ATN signal active, no Identify message is received from the initiator. In this case, the LUN is identified from the CDB and disconnect permission is disabled. The target does not perform any phase retries. The target still responds to a subsequent attention condition. However, the LUN is not considered to be known if a fatal error is detected during the Command phase. That is a Command phase parity error or a fatal message error in response to attention condition during Command phase is handled as a Bus Free error with no sense data. The target also knows the use of linked commands if selected without ATN. The target does not initiate synchronous data transfer negotiation if selected without ATN.

Phase retries and target initiated negotiations may be allowed if a subsequent Identify message is received.

10.15 Multiple Initiator Environment

10.15.1 Initiator Sense Data

Separate sense data is reserved for each initiator. Each initiator's sense data is maintained independent of commands from other initiators.

10.15.2 Initiator Mode Select/Mode Sense Parameters

A single shared copy of the Mode Select/Mode Sense parameters is maintained by the file. This includes both the current and saved parameters.

10.15.3 Initiator Data Transfer Mode Parameter

A separate data transfer mode parameters area is reserved and maintained for each initiator.

10.16 Options Jumper Block

The Drive provides an Options Jumper Block which allows the drive to be customized using jumpers instead of software settings. Refer to Figure 17 on page 23, 6.1.6, “The pin assignment” on page 24, and Figure 27 on page 32 for detailed information on the location and function of these jumpers.

10.17 Contingent allegiance Condition

The contingent allegiance condition shall exist following the return of Check Condition, except Check Condition caused by Invalid LUN. Execution of all queued commands shall be suspended until the contingent allegiance condition is cleared.

The contingent allegiance condition can be cleared by the initiator in one of the following ways:

- By issuing a REQUEST SENSE command to the Target and receiving the sense data. This is most recommended way.
- By issuing any other command to the I_T_x nexus that reported the fault.
- By issuing an Abort message to the I_T_x nexus that reported the fault. This will also abort the current and queued I/O process from that initiator.
- By issuing a Bus Device Reset message to the Target. This will also abort all current and queued I/O processes.
- By generating a RESET condition on the bus. This **MUST** be the last resort.

10.18 Reset

The Reset condition is used to clear all SCSI devices from the bus. This condition takes precedence over all other phases and conditions. After a reset condition is detected and the reset actions completed, the target returns to a 'SCSI bus enabled' state that allows the target to accept SCSI commands.

This device uses the Hard reset option as defined in the SCSI-2 standard.

10.18.1 Reset Sources

There are four sources of resets detected by the target:

Reset Name	Reset Source
Power-On Reset	The 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.
SCSI Bus Reset	This is a reset generated when the SCSI bus control line RST goes active.
SCSI Bus Device Reset Message	This is the reset generated by the SCSI Bus Device Reset Message(0Ch).

10.18.2 Reset Actions

The action taken by the Drive following a reset is dependent on the source of the reset.

10.18.2.1 Power-On reset and Self-Initiated reset

These two reset conditions cause the following to be performed in the order shown.

1. A power-up sequence
2. A start-up sequence is necessary to put the Drive in a ready state

10.18.2.2 SCSI Bus reset and SCSI Bus Device Reset message

These two reset conditions cause the following to be performed.

- If reset goes active while the power-up sequence is in progress, the power-up sequence is started over.
- If the Auto Start pin is grounded and a start-up sequence has not yet completed, a start-up sequence will be re-attempted from the beginning.

Note: The power-up sequence, having already completed, is not rerun.

- If reset occurs while a physical sector is being written, the write operation is disabled after the current physical sector is written. Data is not lost as long as power stays valid until the physical sector being written is completed.

10.19 RAS Characteristics

The file will execute a self test at power on or when a Send Diagnostics command is issued with the self test bit set.

The diagnostics are to assure the correct operation of the file and to verify that the check circuits detect fault conditions.

10.19.1 Power on Diagnostics

At power on time the following tests are executed:

1. Test the microprocessor's:
 - a. Internal Timers.
 - b. Internal RAM.
2. Do a sum check on the microprocessor's external ROM.
3. Test the adapter section of the file as follows:
 - a. Test the registers.
 - b. Test HDC.
4. Do a read/write test on the microprocessor's external RAM.
5. Do a read/write test on sector buffer RAM.
6. Check if the 12 volt line is active.
7. Initialize and check the servo system is functioning correctly.
8. Check the spindle's RPM.
9. Seek to a correct cylinder.

10.19.2 Diagnostics Command

The tests executed as a result of the Send Diagnostics command with the self test bit set to a one differs from the tests executed at power on. The spindle motor must be at the correct speed for the Diagnostics command to be executed. If the motor is not at the correct speed, a NOT READY status will be returned in the sense byte.

The Diagnostics command will execute the following tests:

1. Test the adapter section of the file as follows:
 - a. Test the registers.
 - b. Assure the ECC circuits work correctly.
 - c. Do a read/write test on sector buffer RAM.
2. Check the spindle's RPM.
3. Seek to a correct cylinder.
4. Write the CE cylinder.
5. Read from the CE cylinder.

10.19.3 Diagnostics Fault Reporting

Faults detected before successful completion of the adapter section could prevent the file from responding to a selection.

Faults detected after the successful completion of the adapter 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 command. The Inquiry command will always respond with good status. Detecting a fault during power on will not terminate execution of the diagnostics nor will it terminate the power on process.

Faults detected during a Send Diagnostics command will report a Check Condition as end status.

11.0 SCSI SENSE DATA

11.1 SCSI Sense Data Format

Format of the sense data returned by the file in response to the REQUEST SENSE command.

	BIT							
	7	6	5	4	3	2	1	0
BYTE 0	Valid	Error Code (70h or 71h)						
BYTE 1	RSVD = 0							
BYTE 2	0	ILI	0	Sense Key				
BYTE 3 - 6	(MSB) Information Bytes							(LSB)
BYTE 7	Additional Sense Length							
BYTE 8 -11	(MSB) Product Specific Information							(LSB)
BYTE 12	Additional Sense Code							
BYTE 13	Additional Sense Code Qualifier							
BYTE 14	FRU = 0							
BYTE 15	SKSV	Sense-Key Specific Bits						
BYTE 16 BYTE 17	Sense-Key Specific Bytes							
BYTE 18 -19	Reserved = 0							
BYTE 20 BYTE 21	Unit Error Code							
BYTE 22 BYTE 23	Reserved = 0							
BYTE 24 -27	Product Specific Information							
BYTE 28 -31	Reserved = 0							

Figure 102. Format of Sense Data. Format of the Sense Data returned by the file in response to the REQUEST SENSE command

11.1.1 Sense Data Description

11.1.1.1 Valid Bit

Bit 7 of byte 0

- 0** The Information Bytes are not defined.
- 1** The Information Bytes contain a valid logical block address.

11.1.1.2 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 command is an example of a command that may return a deferred error.

11.1.1.3 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.

0	NO SENSE There is no sense key information to be reported for the logical unit.
1	RECOVERED ERROR The last command completed successfully with some recovery action performed by the file. More detailed information is available in the Additional Sense Code.
2	NOT READY The logical unit addressed cannot be addressed. More detailed information is available in the Additional Sense Code.
3	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.
4	HARDWARE ERROR The file detected a unrecoverable hardware error while performing a command or during a diagnostic test. More detailed information is contained in the Additional Sense Code.
5	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 file might have altered the medium.
6	UNIT ATTENTION Indicates that the file entered in the 'Unit Attention Condition'. (See 10.1.5, "Unit Attention Condition" on page 147.)
7-8	Not used
9	Vendor Specific
A	Not used
B	ABORTED COMMAND The file aborted the command.
C-D	Not Implemented
E	Not Used
F	Reserved

11.1.1.4 Information Bytes

Byte 3 - 6

This field is only valid when Valid Bit is one, and contains the unsigned LBA associated with the sense key.

11.1.1.5 Additional Sense Length

Byte 7

Indicates the remaining number of bytes in the sense data. (It is set to 18h in the file.)

11.1.1.6 Command Specific Information

Byte 8 - 11

Reserved.

Note: The values in this field vary by product.

11.1.1.7 Additional Sense Code/Qualifier

Byte 12 / 13

Key	Code	Qual	Description
0	00	00	No error.
1	01	00	Recovered write error no index
1	03	00	Recovered write error - write fault
1	14	01	Recovered write error Id not found
1	16	00	Recovered write error DAM not found
1	17	00	Recovered read error without ECC applied.
1	17	01	Recovered read error with retries.
1	17	06	Recovered read error without ECC applied. Auto reallocated.
1	17	07	Recovered read error without ECC applied. Recommended reassign. This value can be returned only when ARRE = 0.
1	17	09	Recovered read error without ECC applied. Data re-written. This value can be returned only when ARRE = 1.
1	18	00	Recovered read error with ECC applied. This value can be returned only when ARRE = 0.
1	18	02	Recovered read error with ECC applied. Auto reallocated. This value can be returned only when ARRE = 1.
1	18	05	Recovered read error with ECC applied. Recommended reassign. This value can be returned only when ARRE = 0.
1	18	07	Recovered read error with ECC applied. Data Rewritten. This value can be returned only when ARRE = 1.
1	1C	01	Primary Defect list Not Found. Requested Defect List Format is not supported. Default Lsi Format is returned.(Read Defect Data Only)
1	1C	02	Grown Defect list Not Found. Requested Defect List Format is not supported. Default Lsi Format is returned.(Read Defect Data Only)
1	44	00	Internal target failure
2	04	00	Not ready. Start spindle motor fail.
2	04	01	Not ready. In process of becoming ready.
2	04	02	Not ready. Initializing command required. (Start Unit)
2	04	04	Not ready. Format in progress.
2	31	00	Not ready. Media format corrupt. A format operation was interrupted (power down, reset) prior to completion of a Format Unit command. The Format Unit command should be re-issued and must complete successfully for this error condition to be removed.
2	40	80	Diag Fail - Bring-Up Fail
2	40	85	Diag Fail - RAM Microcode Not Loaded
2	4C	00	Degraded Mode - Self Configuration Fail Configuration/RAM Microcode not loaded

Key	Code	Qual	Description
3	10	00	Medium error. ID CRC error.
3	11	00	Medium error. Unrecovered read error.
3	14	01	Medium error. Record not found.
3	16	00	Medium error. Data synchronization mark error. (DAM error)
3	19	00	Medium error. Defect list error. A defect list error occurs when a data error is detected while reading the manufacturing defect list or while reading or writing the grown defect list.
3	31	01	Medium error. Medium Format Corrupted Reassign Failed
4	01	00	H/W error. No index or sector.
4	02	00	H/W error. No seek complete.
4	03	00	H/W error. Write fault.
4	09	00	H/W error. Track following error.
4	11	00	H/W error. Unrecovered read error in reserved area.
4	31	00	Degrade mode. Format corrupt.
4	32	00	H/W error. No defect spare location available. A no defect spare location available sense code indicates that the Reassign Block command can not proceed the process because all spare sectors have been used, or it will exceed implementation limitation of defect handling of the file.
4	40	80	Degrade Mode. Diagnostic Fail. Configuration sector valid check fail. Reserved area sector valid check fail.
4	40	81	Degrade mode. HDC error.
4	40	82	Degrade mode. HIC error.
4	40	83	Degrade mode. Other LSI error.
4	40	84	Degrade mode. RAM error.
4	40	85	Degrade Mode. RAM Microcode Not Loaded
4	44	00	H/W error. Internal target failure
4	47	00	H/W error. SCSI parity error.
5	1A	00	Illegal request. Parameter list length error. The number of parameters supplied is not equal to the value the expected.
5	20	00	Illegal request. Illegal command operation code. This command is also returned when an unsupported command code is received.
5	21	00	Illegal request. Logical block address out of range.
5	24	00	Illegal request. Invalid field in CDB.
5	25	00	Illegal request. Invalid LUN.
5	26	00	Illegal request. Invalid fields in the parameter list.
6	28	00	Unit attention. Not ready to ready transition.(Format completed)
6	29	00	Unit attention. Power on reset or Bus device reset occurred.
6	2A	01	Unit attention. Mode select parameter changed.

Key	Code	Qual	Description
6	2F	00	Unit attention. Command cleared by another initiator.
6	3F	01	Unit attention. Micro code has been changed.
B	1B	00	Aborted command. Synchronous data transfer error. (Extra ack detected)
B	25	00	Aborted command. Unsupported LUN. The drive supports LUN 0 only.
B	43	00	Aborted command. Message reject error. A message reject error occurs when an inappropriate or unexpected message reject is received from the initiator or the initiator rejects a message twice.
B	45	00	Aborted command. Selection/Reselection failed. A selection/reselection error occurs when the initiator fails to respond to a reselection within 250 milliseconds after the drive starts reselection. The reselection is attempted a second time before setting selection/reselection failed sense code.
B	47	00	Aborted command. SCSI parity error.
B	48	00	Aborted command. Initiator detected error message received. An initiator detected error occurs when the initiator detects an error, sends a message to retry, detects the error again, and sends the retry message a second time. The drive then sets check condition status with Initiator Detected Error.
B	49	00	Aborted command. Inappropriate/illegal message. An inappropriate or illegal message occurs when the initiator sent a message that either is not supported or is not in a logical sequence.
B	4E	00	Aborted command. Overlapped commands attempted.

11.1.1.8 FRU (Field Replaceable Unit)

Byte 14

The FRU (Field Replaceable Unit) field value will always be zero.

11.1.1.9 Sense Key Specific

Byte 15 - 17

The definition of this field is determined by the value of the sense key field.

11.1.1.9.1 Illegal Request Case: Error filed pointer is returned.

	BIT							
	7	6	5	4	3	2	1	0
BYTE 15	SKSV	C/D	Reserved		BPV	Bit Pointer		
BYTE 16	(MSB)		Field Pointer					
BYTE 17							(LSB)	

Figure 103. Field Pointer Bytes

SKSV Sense-key specific valid

C/D Command/Data

0 Indicates that the illegal parameter is in the data parameters sent by the initiator during DATA OUT phase

1 Indicates that the illegal parameter in the command descriptor block.

BPV Bit Pointer Valid

0 Indicates the bit pointer field is not valid.

1 Indicates the bit pointer field is significant.

Bit Pointer 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 point to the most significant bit of the field.

Field Pointer Indicates which bytes of the command descriptor block or of the parameter data was in error. Bytes are numbered starting from zero, as shown in the tables describing the commands and parameters. When a multiple-byte field id in error, the pointer point to the most significant byte of the field.

11.1.1.9.2 Recovered, Hardware or Medium Error Case: Actual Retry Count is reported.

		7	6	5	4	3	2	1	0
		BIT							
BYTE 15	SKSV	Reserved							
BYTE 16	(MSB)	Actual Retry Count							
BYTE 17		(LSB)							

Figure 104. Actual Retry Count Bytes

SKSV Sense-key specific valid
 0 Indicates that Actual Retry Count is not valid.
 1 Indicates that Actual Retry Count is valid.

Actual Retry Count Actual number of retries used in attempting to recover from the error condition.

11.1.1.9.3 Not Ready Case: Progress indication is returned. These fields are only defined for the FORMAT UNIT command with the Immediate bit set to one.

		7	6	5	4	3	2	1	0
		BIT							
BYTE 15	SKSV	Reserved							
BYTE 16	(MSB)	Progress Indication							
BYTE 17		(LSB)							

Figure 105. Format Progress Indication Bytes

SKSV Sense-key specific valid
 0 Indicates that Progress Indication is not valid.
 1 Indicates that Progress Indication is valid.

Progress Indication Indicates a percent complete in which the returned value is the numerator that has 10000h as its denominator.

11.1.1.10 Reserved

Byte 18 - 19

11.1.1.11 Unit Error Code(UEC)

Byte 20 - 21

The UEC 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.

11.1.1.12 Reserved

Byte 22 - 23

11.1.1.13 Physical Error Record

Byte 24 - 27

- ILI = 1 - This field contains zeros.
- ILI = 0 - These bytes contain the physical locations of the error in cylinder, head and sector. Byte 24 and 25 are Cylinder high and cylinder low respectively. Byte 26 is the head number and byte 27 is the sector. If the head and/or sector is undetermined, the value is set to FFh. If the Cylinder value is undetermined, both bytes 24 and 25 are set to FFh. This field is valid with Sense Key 1, 3 and 4 only.

Note: The values in this field vary by product.

11.1.1.14 Reserved

Byte 28 - 31

12.0 Recovery Procedure

The following sections describe the recovery procedure for each of the various types of errors for which recovery is attempted.

12.1 Data Recovery Procedure for Data Field Errors

The following table lists the steps that have been defined for data recovery. These recovery actions are for Sync Byte, Data ECC and No Sector Found errors on read, Verify(2Fh) and Write and Verify(2Eh) commands.

Definitions for following tables.

- Reread - Read with no parameters altered.
- Rewrite - Write with no parameters altered.
- Track offset - Movement of the Head slightly off center. The numbers indicate relative degree of offset 1 = smallest, 3 = largest.
- ECC Correction

An error correction code (ECC) is used to verify and correct data in sector buffer. The 128 bit ECC Reed Solomon code (non-interleaved code of degree 8, operating on 10bit symbols) is used for this code. This code is capable of correcting four 10bit symbols error and detecting up to 8 symbols error (71bit burst). It is guaranteed to correct one 31bit burst error.

- OTF - On The Fly data correction. The hardware calculation for OTF data correction is limited to 4 symbols on adjacent or non-adjacent error. It will correct a 10 bit error and under certain conditions can correct up to 20 bit error.
- OFL - Off Line data correction. Off line data correction is limited to 5 symbols on adjacent or non-adjacent error to reduce the probability of mis-correction. It will correct a 21 bit error and under certain conditions can correct up to 30 bit error.
- Equalizer - The adaptive Channel Equalizer Burst level is set to a value altered by the amount shown.
- Threshold - The adaptive Channel Threshold Level is set to a value altered by the amount shown.
- Data window - The adaptive Channel Data Window Register is set to a value altered by the amount shown.
- Dummy write - Dummy write operation at non data area and read data with no parameters altered.

12.1.1 Read commands

DRP Steps	Action	ECC correction	Miscellaneous
0	Initial Read	OTF	
1-4	Reread	OTF	
5-6	+ 0 track offset	OTF/OFL	
7-8	-0 track offset	OTF/OFL	
9-10	+ 1 track offset	OTF/OFL	
11-12	-1 track offset	OTF/OFL	
13-14	+ 2 track offset	OTF/OFL	
15-16	-2 track offset	OTF/OFL	
17-18	+ 3 track offset	OTF/OFL	
19-20	-3 track offset	OTF/OFL	
21-23	dummy write	OTF	
24	0 Equalizer	OTF	
25	1 Equalizer	OTF	
26	2 Equalizer	OTF	
27	3 Equalizer	OTF	
28	0 Threshold	OTF	
29	1 Threshold	OTF	
30	2 Threshold	OTF	
31	3 Threshold	OTF	
32	0 Data Window	OTF	
33	+ 1 Data Window	OTF	
34	-1 Data Window	OTF	
35	+ 2 Data Window	OTF	
36	-2 Data Window	OTF	
37	+ 3 Data Window	OTF	
38	-3 Data Window	OTF	
39	+ 4 Data Window	OTF	
40	-4 Data Window	OTF	
41-76	Repeat 5-40		
77-112	Repeat 5-40		
13-148	Repeat 5-40		
149	+ 0 track offset	OTF	

Figure 106. Read Recovery Steps.

Note: The track offsets may be restored to zero at the conclusion of the error recovery process. If there are additional sectors to process in a multiple sector operation, offsets remain as is. Offsets are restored if there is an error on the next sector (and entire recovery process begin from the top). Offsets are also restored at the beginning of the next new operation.

Note: When retries are disabled(Read Retry Count = 0 in Mode Page 1), the above listed recovery steps are not taken.

12.1.2 Write Commands

DRP Steps	Action	ECC correction	Miscellaneous
0	Initial Write		
1-80	Rewrite		

Figure 107. Write Recovery Steps.

Note: When retries are disabled(Write Retry Count = 0 in Mode Page 1), the above listed recovery steps are not taken.

12.2 Error Recovery for Non-Data Errors

Errors other than data errors as described above invoke Error Recovery Procedures called ERP. ERP errors fall into 2 classes:

- Position (Servo)
- Spindle Motor

12.2.1 Position (Servo)

DRP Steps	Action	ECC correction	Miscellaneous
0	Initial Seek		
1-20	ReSeek		

Figure 108. Servo Recovery Steps.

12.2.2 Spindle Motor Error

12.2.2.1 POR

DRP Steps	Action	ECC correction	Miscellaneous
0	Initial Spindle start		
1-forever	Spindle restart		

Figure 109. Spindle start Recovery Steps.

Note: Not Ready sense key state will continue until the motor speed reach to the target rotational speed.

12.2.2.2 Start/Stop unit command

A spindle motor error is generated when motor speed can not be controlled. The motor will be stopped and a Not Ready sense key is returned. There is no self initiated recovery taken by the Drive for this error.

12.3 Priority of Error Reporting

Multiple errors can occur during the execution of a command. Since only one error is reported in the sense data, a priority scheme for determining which error to report is used. Following are the rules for reporting sense

- Hard Errors
 - An unrecovered (hard) error will always be reported in place of any soft error that may have occurred during the same operation.
 - If multiple hard error conditions exist, the first hard error detected will be reported.
 - Soft Errors
 - When PER=0 (Mode Select data format page 1),recovered data and non-data errors are not reported.
 - When PER=1 (Mode Select data format page 1),the last error detected will be reported.
- Note:** Error will not be reported even if PER=1 when OTF recovery procedure recovered data.

If an correctable error occurs during idle time,it is reported by the next command processed even though the error may seem inappropriate for that command.

Index

A

ABORT 134
ABORT TAG 136
Attention Condition 139
Automatic rewrite/reallocate 160

B

Back to Back Write 157
BUS DEVICE RESET 136
Byte ordering conventions 45
 Byte ordering conventions 45

C

CLEAR QUEUE TAG 136
Command 41
 FORMAT UNIT 46
 INQUIRY 49
 LOG SELECT 61
 LOG SENSE 61
 MODE SELECT 67
 MODE SENSE 62
 PRE-FETCH 89
 READ 91
 READ BUFFER 98
 READ CAPACITY 92
 READ DEFECT DATA(37) 94
 READ EXTENDED 97
 READ LONG 101
 REASSIGN BLOCKS 102
 RECEIVE DIAGNOSTIC 104
 RELEASE 106
 REQUEST SENSE (03h) 107
 RESERVE 108
 REZERO 110
 SEEK 111
 SEEK EXTENDED 112
 SEND DIAGNOSTIC 113
 START/STOP 116
 SYNCHRONIZE CACHE 117
 TEST UNIT READY 118
 VERIFY 119
 WRITE 120
 WRITE AND VERIFY 122
 WRITE BUFFER 123
 WRITE EXTENDED 121
 WRITE LONG 127
COMMAND COMPLETE 131
Command Processing During Execution of Active I/O
 process 144
Command Processing While Reserved 150
Command queuing 151

Command reordering 154
Concurrent I/O process 156
Contingent Allegiance Condition 167

D

Deferred error 148
Degraded Mode 149
DIAGNOSTIC 113
DISCONNECT 134

F

FORMAT UNIT 46
 Defect List 47

I

I_T_L nexus 107
IDENTIFY 137
Incorrect Initiator Connection 144
INITIATOR DETECTED ERROR 134, 141
INQUIRY 49
Internal Error 148
Invalid LUN 51, 143

L

LBA 45
Link Command 45
 Flag Bits 45
LINKED COMMAND COMPLETE 135
LINKED COMMAND COMPLETE WITH
 FLAG 136
LOG SELECT 61
LOG SENSE 61
Log Sense Data 62
LSB 45
LUN 45

M

MESSAGE PARITY ERROR 135
MESSAGE REJECT 135, 141
MESSAGE SYSTEM 130
 ABORT 134
 ABORT TAG 136
 BUS DEVICE RESET 136
 CLEAR QUEUE TAG 136
 COMMAND COMPLETE 131
 DISCONNECT 134
 IDENTIFY 137
 INITIATOR DETECTED ERROR 134
 LINKED COMMAND COMPLETE 135
 LINKED COMMAND COMPLETE WITH
 FLAG 136

MESSAGE SYSTEM (continued)

MESSAGE PARITY ERROR 135
MESSAGE REJECT 135
NO OPERATION 135
QUEUE TAG MESSAGES(20h, 21h, 22h) 136
RESTORE POINTERS 134
SAVE DATA POINTER 134
SYNCHRONOUS DATA TRANSFER
REQUEST 132

Mode Parameter

Caching Parameters 84
Control Mode Page 86
Disconnect Parameters 78
Format Device Parameters 79
Power Condition Parameters 88
Read-Write Error Recover Parameter 73
Reconnect Parameters 78
Rigid Disk Drive Geometry Parameters 81
Vendor Unique Parameters 70
Verify Error Recover Parameter 82

MODE SELECT 67

Mode select data 69

MODE SENSE 62

Page Control Field. 62
Parameter List 64

MSB 45

Multiple Initiator 166

N

NO OPERATION 135
Non-arbitrating systems 165

O

Options Jumper Block 166

P

Page Format (PF) 67

Parity Error

COMMAND Phase Parity Error 140
DATA OUT Phase Parity Error 140
MESSAGE IN Phase Parity Error 140
MESSAGE OUT Phase Parity Error 140

Power saving mode 158

PRE-FETCH 89

Priority Commands 151

Q

QUEUE TAG MESSAGES(20h, 21h, 22h) 136

R

READ 91

Read Ahead 163

READ BUFFER 98

READ CAPACITY 92
READ DEFECT DATA(37) 94
READ EXTENDED 97
READ LONG 101
REASSIGN BLOCKS 102
RECEIVE DIAGNOSTIC 104
RELEASE 106
REQUEST SENSE 107
RESERVE 108
RESTORE POINTERS 134
REZERO 110
RSVD 45

S

SAVE DATA POINTER 134
Save Page (SP) 67
Saveable Page 67
SCSI Bus Related Error 139
Unexpected Bus Free Error 139

SEEK 111

SEEK EXTENDED 112

Segmented caching 163

Selection without ATN 166

SEND DIAGNOSTIC 113

Sense Data 171

Single Initiator Selection 165

START 116

Status Byte 127

STOP 116

SYNCHRONIZE CACHE 117

SYNCHRONOUS DATA TRANSFER
REQUEST 132

T

TEST UNIT READY 118

Third Party Release 106

Third Party Reserve 108

U

Unit Attention Condition 147

V

VERIFY 119

VU 45

W

WRITE 120

WRITE AND VERIFY 122

WRITE BUFFER 123

Write Cache 157

WRITE EXTENDED 121

WRITE LONG 127



Part Number
39H1055

S39H-1055-03