

Quantum[®]

**Quantum Atlas™ 10K II 9.2/18.4/36.7/73.4 GB
Ultra 160/m S
Product Manual**

June 8, 2000
81-122517-04

Quantum reserves the right to make changes and improvements to products without incurring any obligation to incorporate such changes or improvements in units previously sold or shipped. Quantum reserves the right to make changes to this product manual without incurring any obligation to notify recipients of earlier versions of the product manual.

You can request Quantum publications from your Quantum Sales Representative, or order them directly from Quantum.

Publication Number: 81-122517-04 June 8, 2000

UL / CSA / VDE / TÜV / FCC

UL standard 1954 recognition granted under File No. E78016

CSA standard C22.2-950 certification granted under File No. LR49896

TÜV Rheinland EN 60 950

Tested to FCC Rules for Radiated and Conducted Emissions, Part 15, Sub Part J, for Class-B Equipment.

SERVICE CENTERS

Quantum Service Center
715 Sycamore Avenue
Milpitas, California 95035
Phone: (800) 894-4000
Fax: (408) 894-3218
<http://www.quantum.com>

Quantum Asia-Pacific Pte. Ltd.
50 Tagore Lane #b1-04
Singapore, 2678
Phone: (65) 450-9333
Fax: (65) 452-2544

Quantum Customer Service
Quantum Ireland Ltd.
Finnabair Industrial Park
Dundalk
County Louth, Ireland
Phone: (353) 42-55350
Fax: (353) 45-55355

PATENTS

These products are covered by or licensed under one or more of the following U.S. Patents: 4,419,701; 4,538,193; 4,625,109; 4,639,798; 4,647,769; 4,647,997; 4,661,696; 4,669,004; 4,675,652; 4,703,176; 4,730,321; 4,772,974; 4,783,705; 4,819,153; 4,882,671; 4,920,442; 4,920,434; 4,982,296; 5,005,089; 5,027,241; 5,031,061; 5,084,791; 5,119,254; 5,160,865; 5,170,229; 5,177,771; Other U.S. and Foreign Patents Pending.

© 2000 Quantum Corporation. All rights reserved. Printed in U.S.A.

The Quantum logo is a registered trademark of Quantum Corporation. Quantum and AIRLOCK are trademarks of Quantum Corporation, registered in the U.S.A. and other countries. Capacity for the Extraordinary, Quantum Atlas, DisCache, WriteCache are trademarks of Quantum Corporation. All other brand names or trademarks are the property of their manufacturers.

This product or document is protected by copyright and distributed under licences restricting its use, copying, distributing, and decompilation. No part of this product or document may be reproduced in any form by any means without prior written authorization of Quantum and its licensors, if any.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraphs (c) (1) (ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.227-7013 and FAR 52.227-19.

THIS PUBLICATION IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT.

CONTENTS

REVISION RECORD	xvii
Chapter 1 ABOUT THIS MANUAL	
1.1 AUDIENCE.....	1-1
1.2 DOCUMENT ORGANIZATION.....	1-1
1.3 TERMINOLOGY AND CONVENTIONS.....	1-1
1.4 REFERENCES	1-3
Chapter 2 GENERAL DESCRIPTION	
2.1 PRODUCT OVERVIEW	2-1
2.2 KEY FEATURES	2-1
2.3 STANDARDS AND REGULATIONS.....	2-2
2.4 HARDWARE REQUIREMENTS.....	2-2
Chapter 3 INSTALLATION	
3.1 SAFETY, HANDLING, & ELECTROSTATIC DISCHARGE PROTECTION.....	3-1
3.1.1 Safety Precautions	3-1
3.1.2 Handling.....	3-1
3.1.3 Electrostatic Discharge (ESD) Protection	3-2
3.2 SPACE REQUIREMENTS.....	3-3
3.3 UNPACKING INSTRUCTIONS.....	3-5
3.4 CONFIGURATION JUMPERS AND CONNECTIONS.....	3-6
3.4.1 Jumper Configurations and Connections – Disk Drives with 68-Pin SCSI Connector	3-6
3.4.1.1 TERMPWR (Disk Drives with 68-Pin SCSI Connector)	3-8
3.4.1.2 SCSI ID (Disk Drives with 68-Pin SCSI Connector)	3-8
3.4.1.3 Write Protection (Disk Drives with 68-Pin SCSI Connector).....	3-10
3.4.1.4 Spin Up (Disk Drives with 68-Pin SCSI Connector).....	3-10
3.4.1.5 Enable Narrow Mode (Force 8) (Disk Drives with 68-Pin SCSI Connector).....	3-11
3.4.1.6 Remote Busy and Fault Displays (Disk Drives with 68-Pin SCSI Connector).....	3-11
3.4.2 Jumper Configurations and Connections for SCA-2 Connector Versions	3-12
3.4.2.1 Termination for SCA-2 Connector Versions.....	3-12
3.4.2.2 SCSI ID for SCA-2 Connector Versions.....	3-12
3.4.2.3 Spin Up for SCA-2 Connector Versions	3-13
3.4.2.4 Activity LED for SCA-2 Connector Versions.....	3-14
3.4.2.5 Force Single-Ended Operation for SCA-2 Connector Versions.....	3-15
3.5 CONNECTOR PINOUTS AND WIRING DIAGRAMS	3-16
3.5.1 Atlas 10K II Ultra 160/m SCSI Disk Drive with 68-Pin SCSI Connector.....	3-16
3.5.2 Atlas 10K II Ultra 160/m SCSI Disk Drive with 80-Pin, SCA-2 SCSI Connector	3-18
3.6 DRIVE MOUNTING AND INSTALLATION.....	3-20
3.6.1 Orientation.....	3-20
3.6.2 Mounting Screw Clearance.....	3-21
3.6.3 Shock Clip.....	3-23
3.6.4 Mounting	3-23
3.6.5 Ventilation.....	3-23

3.6.6 Installation in a 3.5-inch Bay	3-24
3.6.7 Installation in a 5.25-inch Bay	3-25
3.6.8 Backplane Installation (SCA Drive)	3-25
3.6.9 Drive Connections	3-26
3.6.9.1 SCSI Bus Connection	3-26
3.6.9.2 Power Connection	3-26
3.6.9.3 Cables	3-27

Chapter 4 SPECIFICATIONS

4.1 SPECIFICATION SUMMARY	4-1
4.2 FORMATTED CAPACITY	4-3
4.3 DATA TRANSFER RATES	4-4
4.4 TIMING SPECIFICATIONS	4-5
4.5 POWER	4-6
4.5.1 Power Sequencing	4-6
4.5.2 Power Reset Limits	4-6
4.5.3 Power Requirements	4-7
4.5.3.1 Power Consumption	4-7
4.5.3.2 DC Voltage Tolerances	4-7
4.5.3.3 Current Requirements	4-8
4.5.3.4 Current Transients	4-9
4.5.3.5 Supply Ripple Voltage and Noise	4-9
4.5.3.6 DC Voltage Monitoring	4-9
4.6 ACOUSTICS	4-10
4.7 MECHANICAL	4-10
4.8 ENVIRONMENTAL CONDITIONS	4-11
4.9 SHOCK AND VIBRATION	4-11
4.10 RELIABILITY	4-12

Chapter 5 SCSI DESCRIPTION

5.1 Overview of the SCSI Command Descriptions	5-3
5.2 Command Descriptor Block	5-5
5.3 Status/Error Reporting	5-9
5.4 Auto Contingent Allegiance Condition and Contingent Allegiance Condition	5-11
5.5 Extended Contingent Allegiance Condition	5-13
5.6 Linked Commands	5-15
5.7 DATA Transfer Command Components	5-17
5.8 SCSI Command Descriptions	5-21
5.9 CHANGE DEFINITION Command (40h)	5-23
5.10 FORMAT UNIT Command (04h)	5-25
5.10.1 Five Forms of FORMAT UNIT Commands	5-27
5.10.2 FORMAT UNIT Parameter List	5-27
5.10.2.1 FORMAT UNIT Defect List Header	5-28
5.10.2.2 FORMAT UNIT Defect Descriptor – Block Format	5-29
5.10.2.3 FORMAT UNIT Defect Descriptor – Physical Sector and Byte From Index Format	5-30
5.10.2.4 FORMAT UNIT Initialization Pattern Descriptor	5-31
5.11 INQUIRY Command (12h)	5-33
5.11.1 Standard Inquiry Data Page	5-35

5.11.2 Vital Product Data Pages	5-37
5.11.2.1 Supported Vital Product Data Pages Page (00h).....	5-37
5.11.2.2 Unit Serial Number Page (80h).....	5-38
5.11.2.3 Implemented Operating Definition Page (81h).....	5-39
5.11.2.4 ASCII Implemented Operating Definition Page (82h).....	5-40
5.11.2.5 Device Identification Page (83h).....	5-41
5.11.2.6 Firmware Revision Page (C0h)	5-43
5.11.2.7 Quantum Manufacturing Information Page (C1h)	5-44
5.11.2.8 Negotiated Rate Page (C4h).....	5-45
5.11.2.9 Command Support Data Pages.....	5-47
5.12 LOG SELECT Command (4Ch).....	5-51
5.13 LOG SENSE Command (4Dh).....	5-53
5.13.1 LOG SENSE Command Descriptor Block	5-55
5.13.2 LOG SENSE Log Pages.....	5-56
5.14 MODE SELECT (6) Command (15h)	5-59
5.14.1 Initiator-Changeable Mode Pages	5-60
5.14.2 Mode Page Types	5-61
5.14.3 Mode Parameter List	5-61
5.14.3.1 Categories of Changeable Pages.....	5-63
5.14.3.2 Unit Attention Condition Page (00h)	5-64
5.14.3.3 Read-Write Error Recovery Page (01h).....	5-65
5.14.3.4 Disconnect-Reconnect Page (02h).....	5-68
5.14.3.5 Verify Error Recovery Page (07h).....	5-71
5.14.3.6 Caching Page (08h)	5-73
5.14.3.7 Control Mode Page (0Ah).....	5-76
5.14.3.8 Notch and Partition Page (0Ch).....	5-79
5.14.3.9 XOR Control Mode Page (10h)	5-82
5.14.3.10 Power Condition Page (1Ah).....	5-84
5.14.3.11 Information Exceptions Control Page (1Ch)	5-87
5.14.3.12 Quantum-Unique Page (39h)	5-91
5.15 MODE SELECT (10) Command (55h).....	5-99
5.16 MODE SENSE (6) Command (1Ah).....	5-103
5.16.1 Read-Only Mode Pages	5-104
5.16.2 Format Device Page (03h).....	5-105
5.16.3 Rigid Disk Geometry Page (04h)	5-107
5.16.4 Vendor-Unique Caching Page (38h).....	5-110
5.17 MODE SENSE (10) Command (5Ah)	5-111
5.18 PERSISTENT RESERVATION IN Command (5Eh).....	5-113
5.19 PERSISTENT RESERVATION OUT Command (5Fh)	5-123
5.20 READ (6) Command (08h)	5-135
5.21 READ (10) Command (28h).....	5-137
5.22 READ BUFFER Command (3Ch).....	5-139
5.23 READ CAPACITY Command (25h).....	5-140
5.24 READ DEFECT DATA (10) Command (37h).....	5-143
5.25 READ DEFECT DATA (12) Command (B7h).....	5-147
5.26 READ LONG Command (3Eh).....	5-151
5.27 READ SKIP MASK Command (E8h).....	5-153
5.28 REASSIGN BLOCKS Command (07h)	5-155

5.29 RECEIVE DIAGNOSTIC RESULTS Command(1Ch)	5-157
5.29.1 Supported Diagnostics Pages Page (00h)	5-159
5.29.2 Translate Address Page (40h)	5-160
5.30 RELEASE (6) Command (17h)	5-163
5.31 RELEASE (10) Command (57h).....	5-165
5.32 REPORT DEVICE IDENTIFIER Command.....	5-167
5.33 REPORT LUNS Command (0Ah).....	5-169
5.34 REQUEST SENSE Command (03h)	5-171
5.34.1 Sense Data Availability.....	5-171
5.34.2 Clearing Sense Data.....	5-172
5.34.3 Status Reporting.....	5-172
5.34.4 Sense Data Format for Error Code 70h and Error Code 71h.....	5-172
5.35 RESERVE (6) Command (16h).....	5-191
5.36 RESERVE (10) Command (56h)	5-193
5.37 REZERO UNIT Command (01h).....	5-197
5.38 SEEK (6) Command (0Bh).....	5-199
5.39 SEEK (10) Command (2Bh)	5-201
5.40 SEND DIAGNOSTIC Command (1Dh).....	5-203
5.40.1 Supported Diagnostic Page List.....	5-204
5.40.2 Translate Address Page.....	5-204
5.41 SET DEVICE IDENTIFIER Command	5-205
5.42 START STOP UNIT Command (1Bh).....	5-207
5.43 SYNCHRONIZE CACHE Command (35h).....	5-209
5.44 TEST UNIT READY Command (00h).....	5-211
5.45 VERIFY Command (2Fh).....	5-213
5.46 WRITE (6) Command (0Ah).....	5-213
5.47 WRITE (10) Command (2Ah)	5-217
5.48 WRITE AND VERIFY Command (2Eh).....	5-219
5.49 WRITE BUFFER Command (3Bh).....	5-221
5.50 WRITE LONG Command (3Fh).....	5-223
5.51 WRITE SAME Command (41h).....	5-225
5.52 WRITE SKIP MASK Command (EAh).....	5-227
5.53 XDREAD (10) Command (52h).....	5-229
5.54 XDWRITE (10) Command (50h)	5-231
5.55 XPWRITE (10) Command (51h).....	5-233

Chapter 6 FEATURE DESCRIPTIONS

6.1 Key Features.....	6-1
6.2 Prefetch	6-2
6.3 Write-Back Caching	6-2
6.4 Zero Latency READ/WRITE.....	6-3
6.5 Disconnect-Reconnect.....	6-3
6.6 Track and Cylinder Skewing	6-3
6.7 Average Access Time.....	6-3
6.8 Embedded Servo System	6-4
6.9 Data Integrity and Security	6-4
6.9.1 Media Error Protection.....	6-4
6.9.2 Transfer Error Protection.....	6-4

6.9.3 Addressing Error Protection	6-4
6.9.4 Data Sector Reallocation Error Protection	6-5
6.9.5 Data Verification	6-5
6.10 Tagged Command Queuing	6-5
6.11 Command Reordering	6-6
6.12 Banded Recording	6-6
6.13 Special Functions.....	6-6
6.14 Power Management.....	6-6
6.15 Diagnostics.....	6-7
6.15.1 Power Up Self Test (POST)	6-7
6.15.2 Periodic Self-Adjustments	6-7
6.15.3 Host Diagnostics.....	6-8
6.15.4 Log Sense	6-8
6.16 Error Recovery	6-8
6.16.1 Seek Error Recovery	6-8
6.16.2 Data Read Errors	6-8
6.16.3 Reallocation of Bad Blocks.....	6-9
6.17 The Ultra 160/m Low Voltage Differential (LVD) SCSI Interface.....	6-9
6.17.1 Diffsens	6-10
6.17.2 Diffsens Switching	6-11
6.17.3 DT and ST Clocking	6-11
6.17.4 Cyclic Redundancy Checking	6-12
6.17.5 Domain Validation.....	6-12
6.18 Hot Plugging/Removal Insertion of SCSI Devices	6-12

Appendix A QUICK REFERENCE

Appendix B SCSI REFERENCE INFORMATION

B.1 SCSI-2 Bus Timing	B-1
B.2 SCSI-3 Bus Timing	B-4
B.3 Signal States.....	B-9
B.3.1 Signal Values	B-9
B.3.2 SCSI ID Bits & Priorities	B-10
B.4 SCSI Signals.....	B-11
B.5 SCSI Bus Phases	B-12
B.5.1 BUS FREE Phase	B-12
B.5.2 ARBITRATION Phase	B-13
B.5.3 SELECTION Phase	B-14
B.5.4 RESELECTION Phase	B-15
B.5.5 Information Transfer Phases	B-17
B.5.5.1 Asynchronous Data Transfer	B-18
B.5.5.2 Synchronous Data Transfer	B-19
B.5.5.3 Wide Data Transfer	B-21
B.5.5.4 Signal Restrictions Between Phases	B-21
B.5.5.5 Message Out – Attention Conditions	B-21
B.6 SCSI Bus Conditions.....	B-23
B.6.1 Attention Condition	B-23
B.6.2 Reset Condition	B-24

B.7 SCSI Message FormatB-26

B.8 Supported SCSI Messages.....	B-29
B.8.1 ABORT TASK Message (0Dh) (<i>ABORT TAG</i>)	B-29
B.8.2 ABORT TASK SET Message (06h) (<i>ABORT</i>)	B-29
B.8.3 CLEAR ACA Message (16h)	B-29
B.8.4 CLEAR TASK SET Message (0Eh)	B-29
B.8.5 CONTINUE TASK SET Message (12h) (<i>CONTINUE I/O PROCESS</i>).....	B-30
B.8.6 DISCONNECT Message (04h)	B-30
B.8.7 IDENTIFY Message (80h - FFh).....	B-30
B.8.8 IGNORE WIDE RESIDUE Message (23h)	B-31
B.8.9 INITIATE RECOVERY Message (0Fh)	B-31
B.8.10 INITIATOR DETECTED ERROR Message (05h).....	B-31
B.8.11 LINKED COMMAND COMPLETE Message (0Ah)	B-33
B.8.12 LINKED COMMAND COMPLETE (WITH FLAG) Message (0Bh).....	B-33
B.8.13 MESSAGE PARITY ERROR Message (09h)	B-33
B.8.14 MESSAGE REJECT Message (07h).....	B-33
B.8.15 NO OPERATION Message (08h).....	B-34
B.8.16 Queue Tag Messages (20h, 21h, or 22h)	B-34
B.8.16.1 HEAD OF QUEUE TAG Message (21h)	B-35
B.8.16.2 ORDERED QUEUE TAG Message (21h)	B-35
B.8.16.3 SIMPLE Message (20h) (<i>SIMPLE QUEUE TAG</i>)	B-35
B.8.17 RELEASE RECOVERY Message (10h).....	B-35
B.8.18 SAVE DATA POINTER Message (02h).....	B-36
B.8.19 SYNCHRONOUS DATA TRANSFER REQUEST Message (01h)	B-37
B.8.20 TARGET RESET Message (0Ch) (<i>BUS DEVICE RESET</i>).....	B-37
B.8.21 TARGET TRANSFER DISABLE Message (13h)	B-37
B.8.22 TASK COMPLETE Message (00h) (<i>COMMAND COMPLETE</i>)	B-38
B.8.23 WIDE DATA TRANSFER REQUEST Message (03h).....	B-38
B.8.24 PARALLEL PROTOCOL REQUEST Message (04h).....	B-40

Appendix C NEGOTIATED RATE INFORMATION PAGE REFERENCE

C.1 TRANSFER PERIOD FACTOR FIELD.....	C-1
C.2 SETTINGS FOR THE QAS, DT, AND IU FIELDS	C-2

INDEX	I-1
--------------------	-----

Figures

3-1	Mechanical Dimensions for the 68-pin Wide SCSI Quantum Atlas 10K II Ultra 160/m Hard Disk Drive (9.1/18.2 GB Models).....	3-3
3-2	Mechanical Dimensions for the 80-pin SCA-2 Quantum Atlas 10K II Ultra 160/m Hard Disk Drive (9.1/18.2 GB Models).....	3-4
3-3	Drive Packing Assembly.....	3-5
3-4	Jumper and Connector Locations, 68-Pin SCSI Connector Version.....	3-7
3-5	Pin Locations on SCA-2 Connector.....	3-13
3-6	Force Single-Ended Operation Jumper Locations.....	3-15
3-7	Connectors on 68-Pin SCSI Connector Versions of Atlas 10K II Ultra 160/m Hard Disk Drives.....	3-16
3-8	SCA-2 Connector on Quantum Atlas 10K II Ultra 160/m SCSI Hard Disk Drive.....	3-18
3-9	Mounting Dimensions for the 68-pin Quantum Atlas 10K II Ultra 160/m SCSI Hard Disk Drive.....	3-20
3-10	Mounting Dimensions for the 80-pin SCA-2 Quantum Atlas 10K II Ultra 160/m SCSI Hard Disk Drive.....	3-21
3-11	Mounting Screw Clearance for the Quantum Atlas 10K II Ultras 160/m SCSI Hard Disk Drive.....	3-22
3-12	Lengthwise Airflow Cooling.....	3-23
3-13	Widthwise Airflow Cooling.....	3-24
3-14	Brackets for 5.25-inch Bay.....	3-25
3-15	Drive Power Connector J2.....	3-26
5-1	Typical (6-Byte) Command Descriptor Block — Data Format.....	5-6
5-2	Command Descriptor Block Control Field — Data Format.....	5-8
5-3	CHANGE DEFINITION Command Descriptor Block — Data Format.....	5-23
5-4	FORMAT UNIT Command Descriptor Block — Data Format.....	5-26
5-5	FORMAT UNIT Parameter List — Data Format.....	5-27
5-6	FORMAT UNIT Defect List Header — Data Format.....	5-28
5-7	FORMAT UNIT Defect Descriptor — Block Format.....	5-29
5-8	FORMAT UNIT Defect Descriptor — Physical Sector and Bytes From Index Format.....	5-30
5-9	FORMAT UNIT Initialization Pattern Descriptor — Data Format.....	5-30
5-10	INQUIRY Command Descriptor Block — Data Format.....	5-33
5-11	Standard Inquiry Data Page — Data Format.....	5-35
5-12	Supported Vital Product Data Pages Page — Data Format.....	5-37
5-13	Unit Serial Number Page — Data Format.....	5-38
5-14	Implemented Operating Definition Page — Data Format.....	5-39
5-15	ASCII Implemented Operating Definition Page — Data Format.....	5-40
5-16	Device Identification Page — Data Format.....	5-41
5-17	Firmware Revision Page — Data Format.....	5-42
5-18	Quantum Manufacturing Information Page — Data Format.....	5-44
5-18A	Negotiated Rate Information Page — Data Format.....	5-45
5-19	Command Support Data Page — Data Format.....	5-47
5-20	LOG SELECT Command Descriptor Block — Data Format.....	5-51
5-21	LOG SENSE Command Descriptor Block — Data Format.....	5-53
5-22	LOG SENSE Log Page Format — Data Format.....	5-56
5-23	Generic Log Parameter — Data Format.....	5-57
5-24	MODE SELECT (6) Command Descriptor Block — Data Format.....	5-59
5-25	Mode Parameter List — Data Format.....	5-61
5-26	Mode Parameter Header (6-Byte) — Data Format.....	5-62
5-27	Mode Parameter Block Descriptor — Data Format.....	5-62
5-28	Unit Attention Condition Page — Data Format.....	5-64
5-29	Read-Write Error Recovery Page — Data Format.....	5-65
5-30	Disconnect-Reconnect Page — Data Format.....	5-68
5-31	Verify Error Recovery Page — Data Format.....	5-71
5-32	Caching Page — Data Format.....	5-73
5-33	Control Mode Page — Data Format.....	5-76
5-34	Notch and Partition Page — Data Format.....	5-79

5-35 XOR Control Mode Page — Data Format5-82

Figures (continued)

5-36	Power Condition Page — Data Format	5-85
5-37	Information Exceptions Control Page — Data Format	5-87
5-38	Quantum-Unique Page — Data Format	5-91
5-39	MODE SELECT (10) Command Descriptor Block - Data Format	5-99
5-40	Mode Parameter Header (10-Byte) — Data Format	5-99
5-41	Mode Parameter Block Descriptor – Data Format	5-100
5-42	MODE SENSE (6) Command Descriptor Block — Data Format	5-103
5-43	Mode Parameter Header (6 Byte) — Data Format	5-104
5-44	Format Device Page — Data Format	5-105
5-45	Rigid Disk Geometry Page — Data Format	5-107
5-46	Vendor-Unique Caching Page — Data Format	5-110
5-47	MODE SENSE (10) Command Descriptor Block - Data Format	5-111
5-48	Mode Parameter Header (10 Byte) — Data Format	5-111
5-49	Mode Parameter Block Descriptor – Data Format	5-112
5-50	PERSISTENT RESERVATION IN Command - Data Format	5-113
5-51	Read Keys Parameters - Data Format	5-115
5-52	Read Reservations Parameters - Data Format	5-116
5-53	PERSISTENT RESERVATION IN Read Reservations Descriptor - Data Format	5-117
5-54	PERSISTENT RESERVATION OUT Command - Data Format	5-123
5-55	PERSISTENT RESERVATION OUT Parameter List - Data Format	5-132
5-56	READ (6) Command - Data Format	5-135
5-57	READ (10) Command — Data Format	5-137
5-58	READ BUFFER Command — Data Format	5-139
5-59	READ CAPACITY Command — Data Format	5-141
5-60	READ CAPACITY Returned Data — Data Format	5-142
5-61	READ DEFECT DATA (10) Command — Data Format	5-143
5-62	Defect Descriptor — Block Format	5-144
5-63	Defect Descriptor — Bytes From Index Format or Physical Sector Format	5-144
5-64	Defect List Header — Data Format	5-145
5-65	READ DEFECT DATA (12) Command — Data Format	5-147
5-66	Defect Descriptor — Block Format	5-148
5-67	Defect Descriptor — Bytes From Index Format or Physical Sector Format	5-148
5-68	Defect List Header — Data Format	5-149
5-69	READ LONG Command Descriptor Block — Data Format	5-151
5-70	READ LONG Command — Returned Data	5-152
5-71	READ SKIP MASK Command Descriptor Block — Data Format	5-153
5-72	REASSIGN BLOCKS Command Descriptor Block — Data Format	5-155
5-73	REASSIGN BLOCKS Defect List Header — Data Format	5-155
5-74	REASSIGN BLOCKS Defect Descriptor — Data Format	5-156
5-75	RECEIVE DIAGNOSTIC RESULTS Command Descriptor Block — Data Format	5-157
5-76	Supported Diagnostic Pages Page — Data Format	5-159
5-77	Translate Address Page — Data Format	5-161
5-78	RELEASE (6) Command Descriptor Block — Data Format	5-163
5-79	RELEASE (10) Command Descriptor Block — Data Format	5-165
5-80	REPORT DEVICE IDENTIFIER Command Descriptor Block — Data Format	5-167
5-81	REPORT DEVICE IDENTIFIER Parameter List — Data Format	5-168
5-82	REPORT LUNS Command Descriptor Block — Data Format	5-169
5-83	LUN Reporting Parameter List — Data Format	5-170
5-84	REQUEST SENSE Command Descriptor Block — Data Format	5-171
5-85	Sense Data Format for Error Code 70h — Data Format	5-173
5-86	ILLEGAL REQUEST Sense Key Field Pointer Bytes — Data Format	5-189
5-87	NOT READY Sense Key - Progress Indication Bytes — Data Format	5-189
5-88	MEDIUM ERROR or RECOVERED ERROR Sense Key - Retry Count — Data Format	5-190

Figures (continued)

5-89	RESERVE (6) Command Descriptor Block — Data Format	5-191
5-90	RESERVE (10) Command Descriptor Block — Data Format.....	5-193
5-91	Extent Descriptors – Data Format	5-195
5-92	Parameter List When LongID and Extent Bits = 1 – Data Format	5-195
5-93	RESERVE (10) ID Only Parameter List – Data Format	5-196
5-94	REZERO UNIT Command Descriptor Block — Data Format	5-197
5-95	SEEK (6) Command Descriptor Block — Data Format	5-199
5-96	SEEK (10) Command Descriptor Block — Data Format.....	5-201
5-97	SEND DIAGNOSTIC Command Descriptor Block — Data Format.....	5-203
5-98	Supported Diagnostic Page List — Data Format	5-204
5-99	Translate Address Page — Data Format.....	5-204
5-100	SET DEVICE IDENTIFIER Command Descriptor Block — Data Format	5-205
5-101	SET DEVICE IDENTIFIER Parameter List — Data Format.....	5-206
5-102	START STOP UNIT Command Descriptor Block — Data Format	5-207
5-103	SYNCHRONIZE CACHE Command Descriptor Block — Data Format	5-209
5-104	TEST UNIT READY Command Descriptor Block — Data Format	5-211
5-105	VERIFY Command Descriptor Block — Data Format.....	5-213
5-106	WRITE (6) Command Descriptor Block — Data Format	5-215
5-107	WRITE (10) Command Descriptor Block — Data Format	5-217
5-108	WRITE AND VERIFY Command Descriptor Block — Data Format	5-219
5-109	WRITE BUFFER Command Descriptor Block — Data Format	5-221
5-110	WRITE LONG Command Descriptor Block — Data Format.....	5-223
5-111	WRITE SAME Command Descriptor Block — Data Format	5-225
5-112	WRITE SKIP MASK Command Descriptor Block — Data Format	5-227
5-113	XDREAD (10) Command Descriptor Block — Data Format	5-229
5-114	XDWRITE (10) Command Descriptor Block — Data Format.....	5-231
5-115	XPWRITE (10) Command Descriptor Block — Data Format	5-233
6-1	ST Clocking and DT Clocking	6-11
B-1	Extended Message – Data Format	B-28
B-2	IDENTIFY Message – Data Format.....	B-31
B-3	Queue Tag Message – Data Format	B-34
B-4	SYNCHRONOUS DATA TRANSFER – Data Format.....	B-36
B-5	WIDE DATA TRANSFER REQUEST Message – Data Format	B-38
B-6	PARALLEL PROTOCOL REQUEST Message – Data Format	B-40

Tables

3-1	SCSI ID Selection on Option Connector (68-Pin SCSI Connector Drives).....	3-9
3-2	Jumper Settings for SCSI ID, 28-Pin Secondary Option Connector	3-10
3-3	SCSI-ID Pin Assignments (SCA-2 Connector Versions of the Disk Drive).....	3-13
3-4	Spin Up on Power On Options.....	3-14
3-5	68-Pin Wide (LVD) SCSI Connector Pin Assignments	3-17
3-6	80-Pin SCA-2 SCSI Connector Pin Assignments.....	3-19
3-7	Connectors and Jumpers.....	3-27
4-1	Specifications.....	4-1
4-2	Formatted Capacity	4-3
4-3	Transfer Rates.....	4-4
4-4	Timing Specifications.....	4-5
4-5	Power Reset Limits	4-6
4-6	Acoustical Characteristics.....	4-10
4-7	Environmental Specifications	4-11
4-8	Shock and Vibration Specifications	4-12
5-1	Supported SCSI Messages.....	5-3
5-2	Command Descriptor Block — Field Descriptions.....	5-6
5-3	Command Descriptor Block Control Field — Field Descriptions.....	5-8
5-4	Status Codes.....	5-10
5-5	DATA-Phase Command Contents.....	5-18
5-6	CHANGE DEFINITION — Field Descriptions.....	5-24
5-7	FORMAT UNIT Command – Field Descriptions.....	5-26
5-8	FORMAT UNIT Command Supported Options	5-27
5-9	FORMAT UNIT Defect List Header — Field Descriptions.....	5-28
5-10	FORMAT UNIT Initialization Pattern Descriptor — Field Descriptions	5-31
5-11	FORMAT UNIT Initialization Pattern Type.....	5-31
5-12	INQUIRY Command Descriptor Block — Field Descriptions.....	5-33
5-13	Standard Inquiry Data Page — Field Descriptions.....	5-36
5-14	Vital Product Data — Page Codes.....	5-37
5-15	Unit Serial Number Page — Field Descriptions.....	5-38
5-16	Implemented Operating Definition Page — Field Descriptions	5-39
5-17	Device Identification Page — Field Descriptions.....	5-42
5-17A	Negotiated Rate Information Page — Field Descriptions.....	5-46
5-18	Command Support Data Page — Field Descriptions	5-48
5-19	Command Support Data Page Command or Operation Codes.....	5-49
5-20	LOG SELECT Command Descriptor Block — Field Descriptions.....	5-51
5-21	Disk Drive Log Pages	5-54
5-22	LOG SENSE Command Descriptor Block — Field Descriptions	5-55
5-23	LOG SENSE Log Page Format – Field Descriptions.....	5-56
5-24	Generic Log Parameter – Field Descriptions.....	5-57
5-25	MODE SELECT (6) Command — Field Descriptions.....	5-59
5-26	Initiator-Changeable Mode Pages	5-60
5-27	Mode Page Types.....	5-61
5-28	Mode Parameter List — Field Descriptions	5-62
5-29	Mode Parameter Header — Field Descriptions.....	5-62
5-30	Mode Parameter Block Descriptor — Field Descriptions.....	5-63
5-31	Categories Of Changeable Pages.....	5-63
5-32	Unit Attention Condition Page — Field Descriptions.....	5-64
5-33	Read-Write Error Recovery Page — Field Descriptions	5-66
5-34	Disconnect-Reconnect Page — Field Descriptions	5-69
5-35	Verify Error Recovery Page — Field Descriptions	5-72
5-36	Caching Page — Field Descriptions.....	5-74

Tables (continued)

5-37	Control Mode Page — Field Descriptions.....	5-77
5-38	Notch and Partition Page — Field Descriptions.....	5-80
5-39	XOR Control Page - Field Descriptions.....	5-83
5-40	Power Condition Page — Field Descriptions.....	5-86
5-41	Information Exceptions Control Page — Field Descriptions.....	5-87
5-42	Codes Used by the MRIE Field.....	5-90
5-43	Quantum Unique Page — Field Descriptions.....	5-92
5-44	Mode Parameter Block Descriptor – Field Descriptions.....	5-100
5-45	MODE SENSE Command — Field Descriptions.....	5-103
5-46	Mode Parameter Header — Field Descriptions.....	5-104
5-47	Read-Only Mode Pages.....	5-104
5-48	Format Device Page — Field Descriptions.....	5-106
5-49	Rigid Disk Geometry Page — Field Description.....	5-108
5-50	Vendor-Unique Caching — Field Description.....	5-110
5-51	Mode Parameter Block Descriptor – Field Descriptions.....	5-112
5-52	PERSISTENT RESERVATION IN Command — Field Descriptions.....	5-113
5-53	Read Keys Parameters — Field Descriptions.....	5-115
5-54	Read Reservations Parameters — Field Descriptions.....	5-116
5-55	PERSISTENT RESERVATION IN Read Reservations Descriptor — Field Descriptions.....	5-117
5-56	Persistent Reservation Type Codes and Their Meanings.....	5-119
5-57	When Do Conflicts Between Existing Reservations and New Reservations Exist?.....	5-122
5-58	PERSISTENT RESERVATION OUT Command — Field Descriptions.....	5-124
5-59	PERSISTENT RESERVATION OUT Command’s Service Action Descriptors.....	5-126
5-60	Persistent Reservation Type Codes and Their Meanings.....	5-130
5-61	PERSISTENT RESERVATION OUT Parameter List – Field Descriptors.....	5-133
5-62	Device Server Interpretation of Service and Scope Value.....	5-134
5-63	READ (6) Command — Field Descriptions.....	5-135
5-64	READ (10) Command — Field Descriptions.....	5-137
5-65	READ BUFFER Command — Field Descriptions.....	5-140
5-66	READ CAPACITY Command — Field Descriptions.....	5-141
5-67	READ DEFECT DATA (10)Command — Field Descriptions.....	5-143
5-68	Defect List Header — Field Descriptions.....	5-145
5-69	READ DEFECT DATA (12)Command — Field Descriptions.....	5-147
5-70	Defect List Header — Field Descriptions.....	5-149
5-71	READ LONG Command Descriptor Block — Field Descriptions.....	5-151
5-72	READ SKIP MASK — Field Description.....	5-153
5-73	REASSIGN BLOCKS Defect List Header — Field Descriptions.....	5-155
5-74	Diagnostic Pages Supported by the Drives.....	5-157

Tables (continued)

5-75	RECEIVE DIAGNOSTIC RESULTS Command Descriptor Block — Field Descriptions	5-158
5-76	Source Descriptor — Field Descriptions	5-162
5-77	RELEASE (6) Command Descriptor Block — Field Descriptions	5-163
5-78	RELEASE (10) Command Descriptor Block — Field Descriptions	5-166
5-79	REPORT DEVICE IDENTIFIER Command Descriptor Block — Field Descriptions	5-167
5-80	REPORT DEVICE IDENTIFIER Parameter List — Field Descriptions	5-168
5-81	REPORT LUNS Command Descriptor Block — Field Descriptions	5-169
5-82	REQUEST SENSE Command Descriptor Block — Field Descriptions	5-171
5-83	Sense Data Fields (Error Code 70h) — Field Descriptions	5-172
5-84	Supported Sense Keys	5-175
5-85	Sense Key Information Field Contents	5-177
5-86	Supported Additional Sense Codes and Qualifiers (in Hex)	5-179
5-87	Sense-Key Specific Field Contents	5-188
5-88	ILLEGAL REQUEST Sense Key Field Pointer Bytes — Field Descriptions	5-189
5-89	NOT READY Sense Key - Progress Indication Bytes — Field Descriptions	5-190
5-90	MEDIUM ERROR OR RECOVERED ERROR Sense Key - Retry Count — Field Descriptions	5-190
5-91	RESERVE (6) Command — Field Descriptions	5-191
5-92	RESERVE (10) Command — Field Descriptions	5-194
5-93	Reservation Types	5-195
5-94	SEND DIAGNOSTIC Command — Field Descriptions	5-203
5-95	SET DEVICE IDENTIFIER Command Descriptor Block — Field Descriptions	5-205
5-96	SET DEVICE IDENTIFIER Parameter List	5-206
5-97	START STOP UNIT Command — Field Descriptions	5-207
5-98	SYNCHRONIZE CACHE Command — Field Descriptions	5-209
5-99	VERIFY Command — Field Descriptions	5-213
5-100	WRITE (6) Command — Field Descriptions	5-215
5-101	WRITE (10) Command — Field Descriptions	5-217
5-102	WRITE AND VERIFY Command — Field Descriptions	5-219
5-103	WRITE BUFFER Command — Field Descriptions	5-222
5-104	WRITE LONG Command — Field Descriptions	5-223
5-105	WRITE SAME Command — Field Descriptions	5-225
5-106	WRITE SKIP MASK Command — Field Descriptions	5-227
5-107	XDREAD (10) Command — Field Descriptions	5-229
5-108	XDWRITE (10) Command — Field Descriptions	5-231
5-109	XPWRITE (10) Command — Field Descriptions	5-233
A-1	SCSI-2/SCSI-3 Equivalent Terminology	A-1
A-2	SCSI-3 Quick Reference – Commands and Messages	A-2
A-3	SCSI-3 Quick Reference – Pages	A-4
A-4	SCSI-3 Quick Reference – Sense Keys	A-5
A-5	SCSI-3 Quick Reference – Status Codes	A-5

Tables (continued)

B-1	SCSI Bus Timing Values – ANSI SCSI-2 Standard.....	B-1
B-2	SCSI-3 Bus Control Timing.....	B-4
B-3	Single Transition (ST) SCSI Bus Data and Information Phase Timing.....	B-5
B-4	Double Transition (DT) SCSI Bus Data and Information Phase Timing	B-6
B-5	Signal Sources.....	B-9
B-6	SCSI ID Bits & Arbitration Priorities	B-10
B-7	SCSI-3 Bus Signal Definitions	B-11
B-8	Information Transfer Phases.....	B-18
B-9	Drive MESSAGE OUT Phase Response.....	B-23
B-10	SCSI Message Format.....	B-26
B-11	Supported SCSI Messages.....	B-27
B-12	Extended Message – Field Descriptions.....	B-28
B-13	Drive’s Response to DISCONNECT Message	B-30
B-14	IDENTIFY Message – Field Descriptions	B-31
B-15	Drive’s Response to INITIATOR DETECTED ERROR.....	B-32
B-16	SYNCHRONOUS DATA TRANSFER REQUEST Message – Field Descriptions	B-37
B-17	WIDE DATA TRANSFER REQUEST Message – Field Descriptions.....	B-39
B-18	PARALLEL PROTOCOL REQUEST Message – Field Descriptions.....	B-41
C-1	Transfer Period Factor Field Values When PARL = 1	C-1
C-2	Transfer Period Factor Field Values When PARL = 0	C-2
C-3	Settings for the QAS, DT, and IU Fields of the Negotiated Rate Information Page (INQUIRY Command)	C-2

REVISION RECORD

This *Revision Record* provides a publication history of this manual. It lists all major revisions and their effective dates. The publication number (part number) is listed in the Revision Level column.

REVISIONS

Revision Level	Date	Summary of Changes
81-122517-01	October 8, 1999	Initial Release Preliminary version of manual.
81-122517-01	December 17, 1999	2 nd release of preliminary manual. Addition of product specifications, changes to INQUIRY command.
81-122517-01	March 3, 2000	3 rd release of preliminary manual. Removal of references to Fibre Channel. Addition of REPORT DEVICE IDENTIFIER and SET DEVICE IDENTIFIER commands. Changes to specifications and SCSI chapter. Added Index.
81-122517-01	April 19, 2000	1 st release of final version of manual. Changes to specifications chapter, SCSI Chapter.
81-122517-02	April 20, 2000	2 nd release of final version of manual. Changes to bit count in explanation of Byte Transfer Length for READ LONG and WRITE LONG commands.
81-122517-03	May 3, 2000	3 rd release of final version of manual. Chapter 4: Changes to specs for number of R/W heads (18.4 GB), numbers supplied for User Track Data Pitch, Servo Sample Interval, Linear Density, Grown Defects from Environmental Change, User Data Blocks/Surface, Interface Transfer Rate, User Data Rate, Average Access to Data, Command Overhead, ECMA seek rate statements. Chapter 5: INQUIRY section, added information for the Negotiated Rate Information Page. Added Appendix C.

(continued)

Revision Record

Revision Level	Date	Summary of Changes
81-122517-04	June 8, 2000	<p>4rd release of final version of manual.</p> <p>Cover: Capacities corrected for 36.7 and 73.4 GB drives.</p> <p>Chapter 4: Changes to specs for formatted capacities (Section 4.2), data transfer rates (Section 4.3), timing specifications (Section 4.4), power consumption (Section 4.5.3.1), current requirements (Section 4.5.3.3), supply ripple voltage and noise (Section 4.5.3.5), acoustics (Section 4.6), and Table 4-8.</p> <p>Chapter 5: Change to <u>ARRE</u> definition (Table 5-33), <u>Reall.Uncorr.Errors</u> definition (Table 5-43). Addition of information for starting AIP following Table 5-43.</p>

This chapter outlines the scope and contents of this manual. It contains information about the intended audience, purpose of the manual, document organization, and document conventions.

1.1 AUDIENCE

This manual was written for original equipment manufacturers (OEMs) that are integrating a Quantum Atlas™ 10K II Ultra 160/m SCSI disk drive into a system or subsystem. Its primary audience is the OEM technical staff that makes disk drive purchase and configuration decisions, and system integrators that are responsible for the SCSI interface. This manual is not intended for end-users and is not a users manual or an installation guide.

1.2 DOCUMENT ORGANIZATION

This product manual is organized into the following chapters:

- Chapter 1 – *About this Manual*
- Chapter 2 – *General Description*
- Chapter 3 – *Installation*
- Chapter 4 – *Specifications*
- Chapter 5 – *SCSI Descriptions*
- Chapter 6 – *Feature Descriptions*

1.3 TERMINOLOGY AND CONVENTIONS

In the Glossary at the back of this manual, you can find definitions for many of the terms used in this manual. In addition, the following abbreviations may be used in this manual:

- ASIC application-specific integrated circuit
- bpi bits per inch
- dB decibels
- dBA decibels, A weighted
- ECC error correcting code
- fci flux changes per inch
- Gbit gigabit
- GB gigabyte
- Hz hertz
- KB kilobytes
- LSB least significant bit

- LVDS low voltage differential SCSI
- mA milliamperes
- MB megabytes (1 MB = 1,000,000 bytes when referring to disk storage and 1,048,576 bytes in all other cases)
- Mbit/S megabits per second
- MB/s megabytes per second
- ms milliseconds
- MSB most significant bit
- mv millivolts
- ns nanoseconds
- SCSI Small Computer System Interface
- tpi tracks per inch
- μ s microseconds
- V volts

The typographical and naming conventions used in this manual are listed below. Conventions that are unique to a specific table appear in the notes that follow that table.

Typographical Conventions:

- Names of Bits: Bit names are presented in initial capitals. An example is the Host Software Reset Bit.
- Commands: Firmware commands are listed as all capitals. An example is MODE SELECT.
- Parameters: Parameters are given as initial capitals when spelled out, and are given as all capitals when abbreviated. Examples are Prefetch Enable (PE) and Cache Enable (CE).
- Hexadecimal Notation: The hexadecimal notation is given as in the example 30h, with the "h" indicating hexadecimal.
- Signal Negation: A signal that is non-active or is in its non-asserted state.
- Messages: A message that is sent from the drive to the host is listed in all capitals. An example is BUS DEVICE RESET.

Naming Conventions:

- Host: In general, the system in which the drive resides is referred to as the host.
- Computer Voice: This refers to any items that the user types at a computer keyboard. These items appear in this manual, if applicable, in 10-point, all capitals, Courier New font. An example is `FORMAT`.

1.4 REFERENCES

For additional information about the SCSI Interface, refer to:

- ANSI Small Computer System Interface-2 (SCSI-2) Specification, ANSI X3T9.2/86-109, Revision 10K II.
- ANSI Small Computer System Interface-3 (SCSI-3) Specification, ANSI X3T10/August, 1994.
- SPI-2 Specification Revision 18.

This chapter summarizes the general functions and key features of the Quantum 10K II Ultra 160/m SCSI family of hard disk drives, and lists applicable standards and regulations.

2.1 PRODUCT OVERVIEW

Quantum Atlas 10K II Ultra 160/m SCSI hard disk drives are part of a family of high performance, 1-inch high (low profile) and 1.6-inch high (half-height) hard disk drives manufactured to meet the highest product quality standards. The disk drives are designed for high-end PC-based workstations, RAID storage systems, digital video applications, database servers, file servers, and other storage systems where high capacity, extraordinary performance, and reliability are critical. These hard disk drives use nonremovable 3.5-inch hard disks and are available with the following SCSI configurations:

- Ultra 160/m, Ultra2, Ultra SCSI 68-pin Wide (16-bit)
- Ultra 160/m, Ultra2, Ultra SCSI 80-pin SCA-2 (16-bit)

There are currently four models in the Atlas 10K II series with capacities of 9.2, 18.4, 36.7, and 73.4 gigabytes (GB).

2.2 KEY FEATURES

Atlas 10K II Ultra 160/m SCSI disk drives include the following features:

- Formatted storage capacity of 9.2 GB (2 disks, 3 heads), 18.4 GB (3 disks, 5 heads), 36.7 GB (5 disks, 10 heads), and 73.4 GB (10 disks, 20 heads). The 9.2, 18.4, and 36.7 GB drives are in the 1.0-inch, low profile form factor; the 73.4 GB drive is in the 1.6-inch, half-height form factor.
- MR heads and banded recording for higher overall bit-densities
- Hardware XOR to automatically detect and correct errors and defects in the data stream for array applications (RAID)
- Self-Monitoring, Analysis, and Reporting Technology (S.M.A.R.T.) Revision 2 support
- ORCA (Optimized Reordering Command Algorithm) and tagged command queuing for greater throughput
- 8 MB Track-oriented segmented cache buffer with prefetch
- Automatic Power Management System with power-saving sub-states
- 10,000 rpm rotational speed yielding 3.0 ms average latency

- 4.7 ms second average random seek time (read) and 7.9 ms average access-to-data time for the 1.0-inch drives; 5.2 ms second average random seek time (read) and 8.4 ms average access-to-data time for the 1.6-inch versions of the drive.
- Embedded servo system for exceptional head positioning accuracy and long life
- “Zero latency” reads and writes reduce data access time
- Full SCSI-3 compliance ¹ (compatible with SCSI-2 and SCSI-1)
- Advanced Ultra 160/m-LVD, Ultra2-LVD and Ultra SE Multimode SCSI with 68-pin wide connector or 80-pin SCA-2 connector
- Superior data integrity:
 - 352-bit Reed-Solomon quad-burst ECC provides fast on-the-fly correction for small errors and robust software-based correction for larger errors
 - EDC internal protection from SCSI bus to media

2.3 STANDARDS AND REGULATIONS

The Quantum Atlas 10K II Ultra 160/m SCSI hard disk drives satisfy the following standards and regulations:

- Underwriters Laboratory (U.L.): Standard 1950. Information technology equipment including business equipment.
- Canadian Standards Association (CSA): Standard C.22.2 No. 950-M89. Information technology equipment including business equipment
- European Standards (TÜV): Standard EN 60 950. Information technology equipment including business equipment
- Federal Communications Commission (FCC): FCC Rules for Radiated and Conducted Emissions, Part 15, Sub Part J, For Class B equipment
- CISPR: CISPR 22 Rules for Radiated and Conducted Emissions, for Class B equipment
- Tested to comply with Australian requirements to carry C-Tick logo

2.4 HARDWARE REQUIREMENTS

The Quantum Atlas 10K II Ultra 160/m SCSI hard disk drives are completely compatible with host computers and controllers that provide a 68-pin Wide or 80-pin SCA-2 interface. A 50-pin to 68-pin adapter is required to use the 68-pin Wide drive in a 50-pin cabling configuration.

¹ The drives are in compliance with the SCSI-3 draft proposed ANSI standards as they exist at the time of product release; SCSI-3 features may not be enabled on drives when shipped.

This chapter explains how to unpack, configure, mount, and connect the Quantum Atlas 10K II Ultra 160/m SCSI hard disk drive prior to operation.

3.1 SAFETY, HANDLING, & ELECTROSTATIC DISCHARGE PROTECTION

3.1.1 Safety Precautions

For your safety, follow all safety procedures described here and in other sections of the manual.

- Remove power from the computer system (or expansion unit) before installing or removing the drive to prevent the possibility of electrical shock or damage to the drive. Unplug the unit containing the drive to provide an added measure of safety.
- Read, understand, and observe all label warnings.

3.1.2 Handling

Damage to the drive can occur as the result of careless handling, vibration, shock, or electrostatic discharge (ESD). Always handle the drive with care to avoid damage to the precision internal components.

CAUTION

A 1/4-inch drop onto a hard surface can damage the drive

Follow these guidelines to avoid damage to the drive:

- Always observe prescribed ESD precautions.
- Keep the drive in its anti-static bag until ready to install.
- Always use a properly fitted wrist strap or other suitable ESD protection when handling the drive.
- Hold drive only by its sides. Do not touch any components on the PCBA.
- Always handle the drive carefully and gently. A drop of 1/4 inch onto a bench or desktop can damage a drive.
- Do not bump, jar, or drop the drive. Use care when transporting the drive.
- Always gently place the drive flat, PCB side down, on an appropriate ESD-protected work surface to avoid the drive being accidentally knocked over.
- Do not pack other materials with the drive in its shielded bag.
- Place the drive in the anti-static bag before placing in shipping container.
- Do not stack objects on the drive.

- Never force the drive or the mounting brackets into the drive bay.
- Do not expose the drive to moisture.
- Do not damage any seals on the drive; doing so may void the warranty.

3.1.3 Electrostatic Discharge (ESD) Protection

Various electrical components within the disk drive are sensitive to static electricity and Electrostatic Discharge (ESD). Even a static buildup or discharge that is too slight to feel can be sufficient to destroy or degrade a component's operation.

To minimize the possibility of ESD-related damage to the drive, we strongly recommend using both a properly installed workstation anti-static mat and a properly installed ESD wrist strap. When correctly installed, these devices reduce the buildup of static electricity which might harm the drive.

Observe the following precautions to avoid ESD-related problems:

- Use a properly installed anti-static pad on your work surface.
- Always use a properly fitted and grounded wrist strap or other suitable ESD protection when handling the drive and observe proper ESD grounding techniques.
- Hold the drive only by its sides. Do not touch any components on the PCBA.
- Leave the drive in its anti-static bag until you are ready to install it in the system.
- Place the drive on a properly grounded anti-static work surface pad when it is out of its protective anti-static bag.
- Do not use the bag as a substitute for the work surface anti-static pad. The outside of the bag may not have the same anti-static properties as the inside. It could actually increase the possibility of ESD problems.
- Do not use any test equipment to check components on the electronics module. There are no user-serviceable components on the drive.

3.2 SPACE REQUIREMENTS

The Quantum Atlas 10K II Ultra 160/m SCSI disk drive is shipped without a faceplate (or bezel) and comes in the following SCSI interface configurations:

- 68-pin Wide SCSI
- 80-pin SCA-2 SCSI

Figures 3-1 and 3-2 show the mechanical dimensions of disk drives of these two interface types.




Figure 3-1 Mechanical Dimensions for the 68-pin Wide SCSI Quantum Atlas 10K II Ultra 160/m Hard Disk Drive (Low Profile Form Factor)




Figure 3-2 Mechanical Dimensions for the 80-pin SCA-2 Quantum Atlas 10K II Ultra 160/m SCSI 9.1/18.2 Disk Drives (Low Profile Form Factor)

Shock Feet

Quantum Atlas 10K II Ultra 160/m SCSI disk drives are outfitted with plastic shock feet on the bottom edge of the base casting, near the corners, beneath the side mounting holes (translucent), and near the corners of the top cover next to the screws (black). The shock feet give an additional level of isolation to prevent the head and disk damage that occasionally may occur during unpacking, staging, and installation. The shock feet attenuate the short-pulse shocks that occur when placing the drive on a hard work surface. If the drive is tested on a hard surface, it should be supported so that its shock feet are not in contact with the hard surface; the disk drive should be supported in the middle, between the shock feet.

NOTE

To provide optimal protection, the shock feet are designed to exceed the form factor when uncompressed.

3.3 UNPACKING INSTRUCTIONS

CAUTION

The maximum limits for physical shock can be exceeded if the disk drive is not handled properly. Special care must be taken not to bump or drop the disk drive.

1. Open the shipping container and remove the packing assembly that contains the disk drive.
2. Remove the disk drive from the packing assembly.

CAUTION

During shipment and handling, the antistatic electrostatic discharge (ESD) bag prevents damage to electronic components due to electrostatic discharge. To avoid accidental damage to the disk drive, do not use a sharp instrument to open the ESD bag. Save the packing materials for possible future use.

3. When you are ready to install the disk drive, remove it from the ESD bag

Figure 3.3 shows the packing assembly for a single Quantum Atlas 10K II Ultra 160/m SCSI hard disk drive. A 12-pack shipping container is available for multiple drive shipments.




Figure 3-3 Drive Packing Assembly

3.4 CONFIGURATION JUMPERS AND CONNECTIONS

This section includes setup and configuration information for Atlas II Ultra 160/m SCSI disk drives. These disk drives include

- The 16-bit multimode Ultra 160/m SCSI-3 wide version with 68-pin SCSI connector,
- The 16-bit multimode Ultra 160/m SCSI-3 version with SCA-2 80-pin connector.

Specific individual settings for each drive type are described in sections 3.4.1 through 3.4.2.

WARNING

Before you begin, review the Safety, ESD, and Handling precautions described at the beginning of this manual to avoid personal injury or damage to equipment.

3.4.1 Jumper Configurations and Connections – Disk Drives with 68-Pin SCSI Connector

This section describes how to configure the jumpers on Atlas 10K II Ultra 160/m disks with 68-pin SCSI interface connectors. The following features are jumper-selectable:

- TERMPWR
- SCSI ID
- Write Protection
- Spin Up
- Remote LED displays (BUSY and FAULT. Note that the FAULT LED illuminates when any diagnostic or functional failure occurs that prevents the drive from performing logical I/O.)

Note that the disk drive does not support on-board SCSI termination.

Figure 3-4 presents an overview of the jumper and connector locations for the wide, 68-pin connector version of the drive.

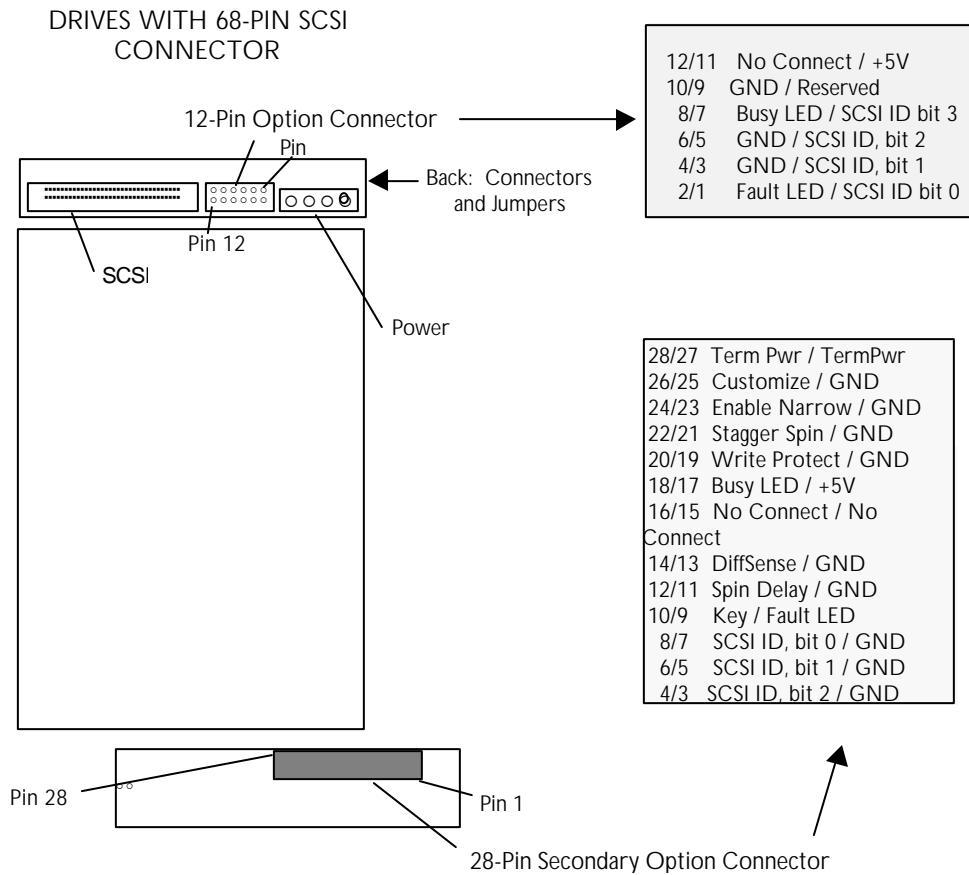


Figure 3-4 Jumper and Connector Locations, 68-Pin SCSI Connector Version

3.4.1.1 TERMPWR (Disk Drives with 68-Pin SCSI Connector)

To ensure that there is a sufficient level of power along the entire SCSI bus, Quantum Corporation recommends that all devices on the SCSI bus supply TERMPWR (if they are capable). SCSI devices are protected by a diode or similar semiconductor to prevent back flow of the terminator power when more than one device supplies this power.

Configure the drive to provide TERMPWR by installing a jumper across pin pair 27 / 28 (Figure 3-4) on the Secondary Option Connector located on the front of the drive.

3.4.1.2 SCSI ID (Disk Drives with 68-Pin SCSI Connector)

Each SCSI device on the bus must have a unique SCSI ID number assigned to it. The drive can be configured for SCSI ID numbers that range from 0 through 15.

Set the SCSI ID for the drive at the 12-pin Option connector or the 28-pin Secondary Option connector. Jumper locations are shown in Figure 3-4. Use Tables 3-1 and 3-2 to assign SCSI IDs using the Option connector or Secondary Option connector pins

NOTE

Refer to your system or SCSI controller documentation for recommendations about assigning SCSI ID numbers for your specific system.

Table 3-1 SCSI ID Selection on Option Connector (68-Pin SCSI Connector Drives)

SCSI ID	Jumper Location – Option Connector			
	Pin Pair 7 / 8	Pin Pair 5 / 6	Pin Pair 3 / 4	Pin Pair 1 / 2
0	0	0	0	0
1	0	0	0	1
2	0	0	1	0
3	0	0	1	1
4	0	1	0	0
5	0	1	0	1
6	0	1	1	0
7	0	1	1	1
8	1	0	0	0
9	1	0	0	1
10	1	0	1	0
11	1	0	1	1
12	1	1	0	0
13	1	1	0	1
14	1	1	1	0
15	1	1	1	1

0 = Jumper Not Installed

1 = Jumper Installed

Configure the drive for remote (external) SCSI ID selection by removing the SCSI ID jumpers (if present) from the referenced SCSI ID pins. Then connect the leads from the external selection switch to the referenced pins. Observe the following guidelines while doing so:

- ID bit 0, at Pin 1, is the Least Significant Bit.
- SCSI ID bits 0, 1, 2, and 3 (pins 1, 3, 5, and 7, respectively) are active LOW signals. That is, the bit is a 1 if the corresponding remote switch is closed to ground.
- Use pins 2, 4, 6 and 8 as the associated ground returns for ID bits 0, 1, 2, and 3, respectively. The ground returns can be used individually, or the ground returns of all four ID bits can be combined at pin 10.

Table 3-2 Jumper Settings for SCSI ID, 28-Pin Secondary Option Connector

SCSI ID	Jumper Location – Secondary Option Connector			
	Pin Pair 7 / 8	Pin Pair 5 / 6	Pin Pair 3 / 4	Pin Pair 1 / 2
0	0	0	0	0
1	1	0	0	0
2	0	1	0	0
3	1	1	0	0
4	0	0	1	0
5	1	0	1	0
6	0	1	1	0
7	1	1	1	0
8	0	0	0	1
9	1	0	0	1
10	0	1	0	1
11	1	1	0	1
12	0	0	1	1
13	1	0	1	1
14	0	1	1	1
15	1	1	1	1

0 = Jumper Not Installed

1 = Jumper Installed

3.4.1.3 Write Protection (Disk Drives with 68-Pin SCSI Connector)

To configure Write Protection for the drive, install a jumper across pin pair 19 / 20 on the Secondary Option Connector (Figure 3-4). To disable Write Protection on the drive, remove the jumper.

3.4.1.4 Spin Up (Disk Drives with 68-Pin SCSI Connector)

Atlas 10K II Ultra 160/m SCSI drives have three Spin Up modes:

- **Option 1: *Spin up immediately when power is applied*** Verify that no jumper is installed across pin pair 11 / 12 (GND / Spin Delay) of the Secondary Option connector (Figure 3-4).
- **Option 2: *Spin up after a predetermined delay following power on***: Install jumpers across pin pair 11 / 12 (GND / Spin Delay) and across pin pair 21 / 22 (GND / Stagger Spin) on the Secondary Option connector (Figure 3-4). Set the delay parameters with the MODE SELECT command, Quantum-Unique Page (39h). The delay is equal to a user-specified multiplier, multiplied by the numerical SCSI ID of the drive. This will give a **staggered spin-up** in multiple-drive installations.
- **Option 3: *Spin up on START STOP UNIT command***: Install a jumper across pin pair 11 / 12 (GND / Spin Delay) and remove any jumper from pin pair 21 / 22 (GND / Stagger Spin), if a jumper is installed there, on the Secondary Option connector (Figure 3-4).

3.4.1.5 Enable Narrow Mode (Force 8) (Disk Drives with 68-Pin SCSI Connector)

Wide Data Transfer (WDTR) negotiations can be limited to 8-bit responses either by setting the Force 8 bit of the Quantum-Unique Page (39h) of the MODE SELECT command (15h) to 1, or by installing a jumper on pin pair 23 / 24 of the Secondary Option Connector (Figure 3-4).

3.4.1.6 Remote Busy and Fault Displays (Disk Drives with 68-Pin SCSI Connector)

Busy and Fault status of the drive can be monitored remotely by connecting a remote (external) Busy and/or remote Fault display LEDs to the appropriate pins on the 12-pin Option Connector or the 28-pin Secondary Option Connector.

Remote Busy LED

On the Option Connector, connect the cathode side of the remote Busy LED to pin 8, Busy LED (Figure 3-4). Connect the anode side of the LED to pin 11, +5V.

On the Secondary Option Connector, connect the cathode side of the remote Busy LED to pin 18, Busy LED (Figure 3-4). Connect the anode side of the LED to pin 17, +5V.

Remote Fault

On the Option Connector, connect the cathode side of the remote Fault LED to pin 2, Fault LED (Figure 3-4). Connect the anode side of the LED to pin 11, +5V.

On the Secondary Option Connector, connect the cathode side of the remote Fault LED to pin 9, Fault LED (Figure 3-4). Connect the anode side of the LED to pin 17, +5V.

3.4.2 Jumper Configurations and Connections for SCA-2 Connector Versions

This section describes jumper settings and connections for Atlas 10K II Ultra 160/m SCSI drives with 80-pin SCA-2 (Single Connector Attachment) connector for the following features:

- SCSI ID
- Spin Up
- Activity LED displays

Use Figure 3-5 to locate the appropriate pins for configuring the drive. Note that Figure 3-5 does not call out each of the 80 pins on the connector, but rather illustrates the layout of the pins.

NOTE

The SCA-2 Connector version of the disk drives do not provide TERMPWR, Active Termination, or Write Protection.

3.4.2.1 Termination for SCA-2 Connector Versions

These versions of Atlas 10K II Ultra 160/m SCSI disk drives cannot be configured to provide bus termination. Therefore, be sure to properly terminate the SCSI bus on which this drive is installed.

NOTE

Refer to your system or SCSI controller documentation regarding any additional recommendations regarding drive placement on the SCSI bus and SCSI bus termination.

3.4.2.2 SCSI ID for SCA-2 Connector Versions

Each SCSI device on the bus must have a unique SCSI ID number assigned to it. The drive can be configured for SCSI ID numbers that range from 0 through 15.

Configure the SCSI ID by providing the proper open or ground signal inputs to the referenced pins of the drive's 80-pin new version SCA-2 connector (Figure 3-5). Refer to Table 3-3 for SCSI ID pin assignments.

NOTE

Refer to your system or SCSI controller documentation for specific recommendations about assigning SCSI ID numbers for your specific system.




Figure 3-5 Pin Locations on SCA-2 Connector

Table 3-3 SCSI ID Pin Assignments (SCA-2 Connector Versions of the Disk Drive)

SCSI ID	Location on SCA Connector			
	ID3 - Pin 80	ID2 - Pin 40	ID1 - Pin 79	ID0 - Pin 39
0	0	0	0	0
1	0	0	0	1
2	0	0	1	0
3	0	0	1	1
4	0	1	0	0
5	0	1	0	1
6	0	1	1	0
7	0	1	1	1
8	1	0	0	0
9	1	0	0	1
10	1	0	1	0
11	1	0	1	1
12	1	1	0	0
13	1	1	0	1
14	1	1	1	0
15	1	1	1	1

0 = Open Circuit, +2.4 V to Vcc +0.5 V

3.4.2.3 Spin Up for SCA-2 Connector Versions

Atlas 10K II drives with SCA-2 SCSI connectors have three Spin Up modes:

- **Option 1** Spin up occurs immediately when power is applied.
- **Option 2** Drive spin up occurs after a predetermined delay following power on. Set the delay parameters with the MODE SELECT Command, Quantum-Unique Page (39h). The delay is equal to a user-specified multiplier multiplied by the numerical SCSI ID of the drive. This will give a staggered spin-up in multiple-drive installations.
- **Option 3** Drive spin up is controlled by the START STOP UNIT command.

Configure the desired spin up option (Table 3-4) by setting the state of the DELAY_SPIN (Pin 38) and STAGGER_SPIN (Pin 78) inputs on the 80-pin SCA-2 connector (Figure 3-5). The states of these signals are set by using either hard-wired connections at the backplane or backplane logic.

Table 3-4 Spin Up on Power On Options

Option	STAGGER_SPIN (Pin 78)	DELAY_SPIN (Pin 38)
Option 1 – Spin Up When Power is Applied	Open	Open
Option 2 – Spin Up After Delay	Ground	Open
Option 3 – Spin Up on START Command	Open	Ground
Reserved	Ground	Ground

3.4.2.4 Activity LED for SCA-2 Connector Versions

The drive provides the output BUSY_OUT signal to power a user-supplied activity LED.

The output indicates the drive is performing a SCSI operation. To use this output, connect a user-supplied LED cathode to the BUSY_OUT connection, pin 77 on the SCA Connector (Figure 3-5). The LED anode must be attached to the proper +5 VDC supply through an appropriate current-limiting resistor.

3.4.2.5 Force Single-Ended Operation for SCA-2 Connector Versions


A jumper is provided to force the disk drive to operate as a single-ended device. See the explanation of the LVD SCSI interface in Chapter 7 for details.

To force the disk drive to operate as a single-ended device, install the low profile jumper across pin pair 1 / 3 (Figure 3-6). For LVD operation and monitoring of DIFFSENS signal, place the jumper across pin pair 2 / 4.

Figure 3-6 Force Single-Ended Operation Jumper Locations

Force single-ended jumper shown installed across Pin Pair 1 / 3. This forces the disk to operate as a single-ended device and disables its ability to monitor DIFFSENS.

Remove the jumper or park it across Pin Pair 2 / 4 to enable the disk drive's DIFFSENS monitoring feature, required for LVD SCSI.

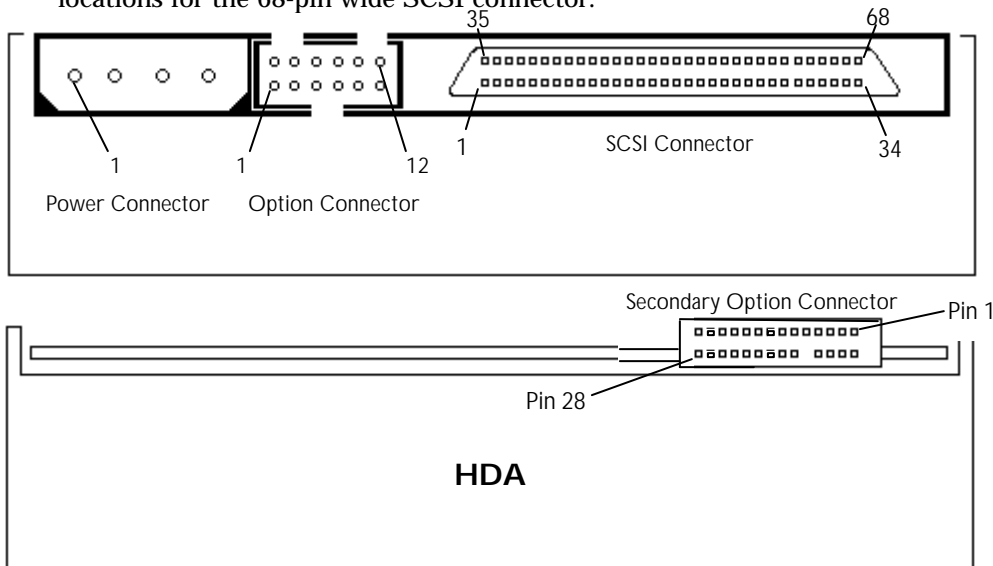


3.5 CONNECTOR PINOUTS AND WIRING DIAGRAMS

Atlas 10K II Ultra 160/m SCSI disk drive interfaces include a 16-bit wide SCSI version and a SCA-2 connector version. Connector pinout information for the drive and signal connection diagrams are found in this section.

3.5.1 Atlas 10K II Ultra 160/m SCSI Disk Drive with 68-Pin SCSI Connector

Figure 3-7 shows the locations of the Power, SCSI, Option, and Secondary Option Connectors on the 68-pin SCSI connector version of the drive. Figure 3-7 also provides the locations and names of the signals connected to the pins at the Option and Secondary Option Connectors on the drive. Table 3-5 provides signal names/pin locations for the 68-pin wide SCSI connector.



Option Connector Signals	
Pin - Signal	Pin - Signal
1 - SCSI ID, bit 0	2 - Fault LED
3 - SCSI ID, bit 1	4 - GND
5 - SCSI ID, bit 2	6 - GND
7 - SCSI ID, bit 3	8 - Busy LED
9 - Reserved	10 - GND
11 - +5 V	12 - No Connect

Secondary Option Connector Signals	
Pin - Signal	Pin - Signal
1 - GND	2 - SCSI ID, bit 3
3 - GND	4 - SCSI ID, bit 2
5 - GND	6 - SCSI ID, bit 1
7 - GND	8 - SCSI ID, bit 0
9 - Fault LED	10 - key
11 - GND	12 - Spin Delay
13 - GND	14 - DiffSense
15 - No Connect	16 - No Connect
17 - +5V	18 - Busy LED
19 - GND	20 - Write Protect
21 - GND	22 - Stagger Spin
23 - GND	24 - Enable Narrow
25 - GND	26 - Customize
27 - TermPwr	28 - TermPwr

Figure 3-7 Connectors on 68-Pin SCSI Connector Versions of Atlas 10K II Ultra 160/m Hard Disk Drives

Table 3-5 68-Pin Wide (LVD) SCSI Connector Pin Assignments

Signal Name	Pin Number	Pin Number	Signal Name
+DB(12)	1	35	-DB(12)
+DB(13)	2	36	-DB(13)
+DB(14)	3	37	-DB(14)
+DB(15)	4	38	-DB(15)
+DB(P1)	5	39	-DB(P1)
+DB(0)	6	40	-DB(0)
+DB(1)	7	41	-DB(1)
+DB(2)	8	42	-DB(2)
+DB(3)	9	43	-DB(3)
+DB(4)	10	44	-DB(4)
+DB(5)	11	45	-DB(5)
+DB(6)	12	46	-DB(6)
+DB(7)	13	47	-DB(7)
+DB(P)	14	48	-DB(P)
GROUND	15	49	GROUND
DIFFSENS	16	50	GROUND
TERMPWR	17	51	TERMPWR
TERMPWR	18	52	TERMPWR
RESERVED	19	53	RESERVED
GROUND	20	54	GROUND
+ATN	21	55	-ATN
GROUND	22	56	GROUND
+BSY	23	57	-BSY
+ACK	24	58	-ACK
+RST	25	59	-RST
+MSG	26	60	-MSG
+SEL	27	61	-SEL
+C/D	28	62	-C/D
+REQ	29	63	-REQ
+I/O	30	64	-I/O
+DB(8)	31	65	-DB(8)
+DB(9)	32	66	-DB(9)
+DB(10)	33	67	-DB(10)
+DB(11)	34	68	-DB(11)

NOTE

A negative sign (-) before a signal indicates an active-low signal.

3.5.2 Atlas 10K II Ultra 160/m SCSI Disk Drive With 80-Pin, SCA-2 SCSI Connector

Figure 3-8 shows the location of the SCSI Connector on this version of the drive. Table 3-6 provides signal names/pin locations for the 80-pin SCA-2 SCSI connector.




Figure 3-8 SCA-2 Connector on Quantum Atlas 10K II Ultra 160/m SCSI Hard Disk Drive

Table 3-6 80-Pin SCA-2 SCSI Connector Pin Assignments

Signal Name	Pin Number	Pin Number	Signal Name
+12 V	1	41	12V GROUND
+12 V	2	42	12V GROUND
+12 V	3	43	12V GROUND
+12 V	4	44	MATED 1
NO CONNECT	5	45	NO CONNECT
NO CONNECT	6	46	SCSI_DIFFSENS
-DATA_M(11)	7	47	DATA_P(11)
-DATA_M (10)	8	48	DATA_P(10)
-DATA_M B(9)	9	49	DATA_P(9)
-DATA_M(8)	10	50	DATA_P(8)
-SCSI_I/O	11	51	SCSI_I/O_P
-SCSI_REQ	12	52	SCSI_REQ_P
-SCSI_C/D	13	53	SCSI_C/D_P
-SCSI_SEL	14	54	SCSI_SEL_P
-SCSI_MSG	15	55	SCSI_MSG_P
-SCSI_RST	16	56	SCSI_RST_P
-SCSI_ACK	17	57	SCSI_ACK_P
-SCSI_BSY	18	58	SCSI_BSY_P
-SCSI_ATN	19	59	SCSI_ATN_P
-SCSI_PARITY_M(0)	20	60	SCSI_PARITY_P(0)
-DATA_M(7)	21	61	DATA_P(7)
-DATA_M(6)	22	62	DATA_P(6)
-DATA_M(5)	23	63	DATA_P(5)
-DATA_M(4)	24	64	DATA_P(4)
-DATA_M(3)	25	65	DATA_P(3)
-DATA_M(2)	26	66	DATA_P(2)
-DATA_M(1)	27	67	DATA_P(1)
-DATA_M(0)	28	68	DATA_P(0)
-SCSI_PARITY_M(1)	29	69	SCSI_PARITY_P(1)
-DATA_M(15)	30	70	DATA_P(15)
-DATA_M(14)	31	71	DATA_P(14)
-DATA_M(13)	32	72	DATA_P(13)
-DATA_M(12)	33	73	DATA_P(12)
+5V	34	74	GROUND (MATED 2)
+5V	35	75	5V GROUND
+5V	36	76	5V GROUND
RESERVED	37	77	ACTIVE LED OUT
-DELAY_SPIN	38	78	-STAGGER_SPIN
-SCSI ID (0)	39	79	-SCSI ID (1)
-SCSI ID (2)	40	80	-SCSI ID (3)

NOTE

A negative sign (-) before a signal indicates an active-low signal.

3.6 DRIVE MOUNTING AND INSTALLATION

This section provides information for installing the disk drive and for connecting the SCSI and power cables as applicable.

WARNING

Before you begin, review the Safety, ESD, and Handling precautions described in the beginning of this manual to avoid personal injury or damage to the drive.

3.6.1 Orientation

The drive can be mounted in any position. Figures 3-9 and 3-10 show the drive base assembly, with mounting locations and dimensions indicated.




Figure 3-9 Mounting Dimensions for the 68-pin Quantum Atlas 10K II Ultra 160/m SCSI Hard Disk Drive




Figure 3-10 Mounting Dimensions for the 80-pin SCA-2 Quantum Atlas 10K II Ultra 160/m SCSI Hard Disk Drive

3.6.2 Mounting Screw Clearance

CAUTION

The printed circuit board assembly (PCBA) is very close to the mounting holes. Do not exceed the specified length for the mounting screws. The specified screw length allows full use of the mounting hole threads while avoiding damaging or placing unwanted stress on the PCBA.

Figure 3-11 specifies the minimum clearance between the PCBA and the screws in the mounting holes. The maximum torque applied to the screws must not exceed 14 inch-pounds (1.57 N-m). A maximum screw length of 0.22 inches (5.59 mm) may be used in the side mounting locations when a bracket of 0.04 inches (1.016 mm) minimum thickness is used.

The PCBA is very close to the mounting holes. Figure 3-11 specifies the clearance between the screws in the mounting holes and the PCBA. Do not use mounting screws longer than the maximum lengths specified in Figure 3-11. The specified screw length allows full use of the mounting hold threads while avoiding damage or stress on the PCBA. Clearance from the disk drive to any other surface (except mounting surfaces) must be a minimum of 0.05 inches (1.25 mm).




Figure 3-11 Mounting Screw Clearance for the Quantum Atlas 10K II Ultra 160/m SCSI Hard Disk Drive

3.6.3 Shock Clip

The shock clip provides an extra level of isolation from the occurrence of a short duration, non-operational shock, particularly when the disk drive is mounted on a rigid platform.

CAUTION

To prevent the possibility of a cross-threaded screw, it is recommended that the screws associated with the fixed side mounting locations be started before the screw associated with the shock clip. Install all of the applicable screws before the final torquing process.

Shock performance is optimal when either four bottom mount locations or four outer side screw locations are used.

3.6.4 Mounting

For the best results during performance benchmark testing, it is recommended that the disk drive be mounted firmly within a system or a fixture, rather than sitting unconstrained on a tabletop or work surface. If it is necessary to perform testing on an unconstrained disk drive, place the disk drive on a flat, smooth, semi-cushioned surface (similar to a mousepad). This avoids any performance degradations due to possibly higher incidence of recovered errors. The disk drive should not be operated on a hard surface.

3.6.5 Ventilation

The Quantum Atlas 10K II Ultra 160/m SCSI hard disk drive operates without a cooling fan, provided the ambient air temperature does not exceed 131°F (55°C). Figures 3-12 and 3-13 show possible configurations for airflow required for adequate cooling. Clearance from the drive to any other surface above and below the disk drive must be a minimum of 1/16-inch (1.25 mm).




Figure 3-12 Lengthwise Airflow Cooling




Figure 3-13 Widthwise Airflow Cooling

3.6.6 Installation in a 3.5-inch Bay

CAUTION

Screw-length must not exceed 0.25-inch (6.25 millimeters) or the drive may be damaged.

The drive base contains threaded holes on its sides and bottom to mount the drive (see Figure 3-9 or Figure 3-10). The holes accept #6-32 screws.

NOTE

Mounting brackets may be used when installing the drive in a 5.25-inch form factor bay. Brackets supplied by Quantum Corporation to mount the drive are recommended; otherwise, the drive's isolation characteristics may be affected. A Bracket kit (part no. 70-30695-01) is available from the Quantum Corp.

To mount the drive:

1. Position the completed drive in the drive bay and choose the drive base mounting holes that match the bay mounting locations.
2. Insert and secure the screws at the selected locations.

3.6.7 Installation in a 5.25-inch Bay

CAUTION

Screws used to mount the drive must not exceed 0.25-inch (6.25 millimeters) in length, or the drive may be damaged.

Figure 3-14 shows the brackets representative of those used to mount the drive in a 5.25-inch bay. A bracket kit including four screws is available from Quantum Corporation. The part number of the kit is 70-30695-01.




Figure 3-14 Brackets for 5.25-inch Bay

Attach the brackets to the drive using four #6-32 screws (see Figure 3-11).

Mount the drive in the bay by performing the following steps:

1. Position the completed unit (drive and brackets) in the bay.
2. Secure the unit with screws at the selected mounting hole locations.

3.6.8 Backplane Installation (SCA Drive)

CAUTION

Screws used to mount the drive must not exceed 0.25-inch (6.25 millimeters) in length, or the drive may be damaged.

Because of the variety of mounting possibilities, these instructions are general in nature. They should be used only as a guide for mounting the drive in your particular system.

Mount the drive as follows:

1. Position the drive in the bay.
2. Align the RAID connector on the drive to the backplane connector in the system.

3. Insert the RAID connector of the drive into the system backplane connector, being careful not to damage any connector pins.
4. Secure the drive in the RAID system with #6-32x1/4 UNC screws. Use the side mount or bottom mount holes, depending upon enclosure mounting hole requirements.

NOTE

To ensure proper operation of the drive, run any required software after the installation is completed.

3.6.9 Drive Connections

The drive connects to both the SCSI and power buses.

3.6.9.1 SCSI Bus Connection

Align the system SCSI bus cable connector to J1 on the drive. Connect the cable to J1 carefully to avoid bending or damaging the connector pins. Drives with SCA connector may plug directly into a backplane connector

The connectors used on the drive vary with the drive type:

- Wide Ultra3 SCSI-3 drives with 68-pin SCSI Connector – Combination Connector containing 68-pin SCSI connector, 4-pin Power Connector, and 12-pin Option Connector (Figure 3-7).
- Wide Ultra3 SCSI-3 drives with 80-pin SCA-2 connector providing SCSI, Power, and configuration/option inputs (Figure 3-8).

3.6.9.2 Power Connection

Align the system power cable connector to the drive power connector J2 (Figure 3-15). Connect the system power cable to J2 carefully to avoid bending or damaging the connector pins. Note that drives with SCA-2 connector do not have this separate J2 Power Connector.




Figure 3-15 Drive Power Connector J2

3.6.9.3 Cables

Furnish the appropriate cables and connectors to match the drive connectors listed in Table 3-7.

Table 3-7 Connectors and Jumpers

Connector/Jumper	Quantum (or Vendor)	Part Number
16-bit Wide SCSI (68-Pin)		
68-pin SCSI connector (includes a 68-pin SCSI connector, a 4-pin power connector, and a 2 mm- pitch, 12-pin (2x6) option connector)	Quantum	23-409326-01
	Molex	87360-0009
23-pin secondary option connector [2 mm-pitch 2X14 keyed connector]	Quantum	TBD
	JAE	TBD
Single Connector Attachment SCSI (80-Pin, SCA-2)		
RAID connector (Champ 0.050 inch Series I 80-position plug)	Quantum AMP	23-416824-80 1123283-9
Shunt Jumper (All drives)		
2 mm-pitch shunt for configuration of drive characteristics	Quantum Augat Berg	23-400479-01 CN-AX027C5A02 86730-001

Chapter 4 SPECIFICATIONS

This chapter provides a detailed description of the physical, electrical, and environmental characteristics of the Quantum Atlas 10K II 160/m SCSI hard disk drives.

4.1 SPECIFICATION SUMMARY

Table 4-1 Specifications

DESCRIPTION	QUANTUM ATLAS 10K II 9.1 GB	QUANTUM ATLAS 10K II 18.4 GB	QUANTUM ATLAS 10K II 36.7 GB	QUANTUM ATLAS 10K II 73.4 GB
Formatted Capacity	9.10 GB	18.4 GB	36.7 GB	73.4 GB
Nominal Rotational Speed (rpm)	10,000	10,000	10,000	10,000
Number of Disks	2	3	5	10
Number of R/W Heads	3	5	10	20
Maximum Areal Density (Gb/square inch)	7.7	7.7	7.7	7.7
Data Organization:				
User Data Tracks per Surface	17337	17337	17337	17337
Data Bands per Surface	24	24	24	24
User Data Track Pitch (at 14,200 tpi)	22500	22500	22500	22500
Data Blocks/Track	301 – 528 (ID – OD)	301 – 528 (ID – OD)	301 – 528 (ID – OD)	301 – 528 (ID – OD)
Data Cylinders	17337	17337	17337	17337
Servo Sectors:				
Servo Spokes/Track	132 (all tracks)	132 (all tracks)	132 (all tracks)	132 (all tracks)
Servo Sample Interval	45.45 μ sec	45.45 μ sec	45.45 μ sec	45.45 μ sec
Bytes Per Sector	512 to 524	512 to 524	512 to 524	512 to 524
Recording:				
Recording Technology	ID-less Split Sector	ID-less Split Sector	ID-less Split Sector	ID-less Split Sector
Linear Density (Kfci)	ID = 341.6 OD = 306.1	ID = 341.6 OD = 306.1	ID = 341.6 OD = 306.1	ID = 341.6 OD = 306.1
Encoding Method	50/52 RLL PRML	50/52 RLL PRML	50/52 RLL PRML	50/52 RLL PRML

Table 4-1 Specifications (continued)

DESCRIPTION	QUANTUM ATLAS 10K II 9.1 GB	QUANTUM ATLAS 10K II 18.2 GB	QUANTUM ATLAS 10K II 36.4 GB	QUANTUM ATLAS 10K II 72.8 GB
Buffer Size	8 MB	8 MB	8 MB	8 MB
RELIABILITY:				
Annualized Failure Rate (AFR) ¹	No greater than 1.5%	No greater than 1.5%	No greater than 1.5%	No greater than 1.5%
Seek Error Rate	10 in 10 ⁶	10 in 10 ⁶	10 in 10 ⁶	10 in 10 ⁶
Recoverable Error Rate	10 in 10 ¹²	10 in 10 ¹²	10 in 10 ¹²	10 in 10 ¹²
Unrecoverable Error Rate	10 in 10 ¹⁴	10 in 10 ¹⁴	10 in 10 ¹⁴	10 in 10 ¹⁴
Grown Defects from Environmental Change (Maximum) ²	18	36	72	144
Minimum Contact Start/Stop Cycles @ 25°C (77°F)	40,000	40,000	40,000	40,000
Auto Head-Park Method	AirLock® – with magnetic bias	AirLock® – with magnetic bias	AirLock® – with magnetic bias	AirLock® – with magnetic bias

NOTES

- ¹ Projected AFR is from a predicted theoretical AFR for the Atlas 10K II family of products that is based on design (i.e., Bellcore) and historical data and does not include process variance, returns with no trouble found, or handling and excessive shock failures. Historically, the field AFR, which returns all returns regardless of cause, has been 50 – 60% of the projected AFR. The development of an operational AFR methodology and derating curve, in line with the IDEMA AFR standardization effort, is underway. For more information, see the HDD Reliability White Paper on Quantum Corporation’s World Wide Web site at www.quantum.com.
- ² With AWRE and ARRE set to 1, the drive will add detected bad blocks to the Grown Defect List and reallocate the data. A small number of grown defects can occur, typically during the first 48 hours of I/O activity, as a result of significant environmental change. This change includes specification extremes (altitude, voltage, temperature, shock, vibration, etc.) not encountered during the manufacturing test process. Environmental extremes and shocks encountered during shipping and handling may also lead to grown defects.

4.2 FORMATTED CAPACITY

The Quantum Atlas 10K II Ultra 160/m SCSI hard disk drives receive a low-level format at the manufacturing facility. This formatting creates the actual tracks and sectors on the disk drive. Table 4-2 shows the storage capacities of the three different disk drives that result from this process. Formatting done at the user level, for operation with DOS, UNIX, or other operating systems, will result in less capacity than the physical capacity shown.

Table 4-2 Formatted Capacity

DESCRIPTION	QUANTUM ATLAS 10K II 9.2 GB	QUANTUM ATLAS 10K II 18.4 GB	QUANTUM ATLAS 10K II 36.7 GB	QUANTUM ATLAS 10K II 73.4 GB
Formatted Capacity	9.2 GB	18.4 GB	36.7 GB	73.4 GB
User Data Blocks/Surface	17,938,986	35,860,910	71,721,820	143,443,640

4.3 DATA TRANSFER RATES

Table 4-3 shows the data transfer rates for the Quantum Atlas 10K II Ultra 160/m SCSI disk drives. Specifications are for both the 68-pin Wide and SCA-2 connector versions of the disk drive.

Table 4-3 Transfer Rates

DESCRIPTION	QUANTUM ATLAS 10K II 9.1 GB	QUANTUM ATLAS 10K II 18.2 GB	QUANTUM ATLAS 10K II 36.7 GB	QUANTUM ATLAS 10K II 73.4 GB
Interface Transfer Rate (MB/second)	160	160	160	160
Media Transfer Rate				
Channel Rate	38.2 – 26.09 MB/second	38.2 – 26.09 MB/second	38.2 – 26.09 MB/second	38.2 – 26.09 MB/second
User Data Rate	42.0 – 31.0 MB/second	42.0 – 26.0 MB/second	42.0 – 19.5 MB/second	42.0 – 19.5 MB/second
Data Streaming Rate (min/max, MB per sec.)	29/40	24/40	24/40	24/40
Weighed Ave. Streaming Rate (MB per sec.)	22.77	22.77	22.77	22.77

4.4 TIMING SPECIFICATIONS

Table 4-4 illustrates the timing specifications of the Quantum Atlas 10K II Ultra 160/m SCSI disk drives.

Table 4-4 Timing Specifications

DESCRIPTION	QUANTUM ATLAS 10K II 9.2 GB	QUANTUM ATLAS 10K II 18.4 GB	QUANTUM ATLAS 10K II 36.7 GB	QUANTUM ATLAS 10K II 73.4 GB
Seek Times (All Nominal):				
Single Track Seek	.6 ms	.6 ms	.6 ms	.6 ms
Average Seek (READ)	4.7 ms	4.7 ms	4.7 ms	5.2 ms
Average Seek (WRITE)	4.7 ms	4.7 ms	4.7 ms	5.2 ms
Average Access to Data	7.9 ms	7.9 ms	7.9 ms	8.4 ms
Average Head Switch Time	0.6 ms	0.6 ms	0.6 ms	0.6 ms
Full Stroke Seek (Max. Seek)	12 ms	12 ms	12 ms	13 ms
Average Rotational Latency	3 ms	3 ms	3 ms	3 ms
Command Overhead, Seek	200 us	200 us	200 us	200 us
Start/Stop Times:				
Time to Data Available ¹ (from SCSI START command)	Nominal: 15 seconds	Nominal: 15 seconds.	Nominal: 15 seconds	Nominal: 17 seconds
Time to Data Available ² (from power up)	Nominal: 25 seconds	Nominal: 25 seconds	Nominal: 25 seconds	Nominal: 27 seconds
Spin Down Time	12 seconds	12 seconds	12 seconds	14 seconds

NOTES FOR TABLE 4-4

- ¹ Time to Data Available from SCSI start command assumes that the SPIN DELAY jumper is installed and power has been supplied to the disk drive with the PCBA ready. Measurement is from the issue of a SCSI START command until a good test unit status is received from the disk drive.
- ² Time to Data Available from power up assumes that no SPIN DELAY jumper is installed. Measurement is from the time power is applied until a good test unit status is received from the disk drive.

4.5 POWER

The Quantum Atlas 10K II Ultra 160/m SCSI hard disk drive operates from two supply voltages:

+ 12V \pm 5%

+ 5V \pm 5%

The allowable ripple and noise (60Hz - 100MHz) is 150 mV peak-to-peak, for the + 12V supply and 75 mV peak-to-peak, for the + 5V supply.

4.5.1 Power Sequencing

Power may be applied in any order or manner, or open either the power or power return line with no loss of data or damage to the disk drive. However, data may be lost in the sector being written at the time of a power loss. The disk drive can withstand transient voltages of + 10% to -100% from nominal while powering up or down.

4.5.2 Power Reset Limits

When powering up, the disk drive remains reset until both V_{HT} reset limits in Table 4-5 are exceeded. When powering down, the drive becomes reset when either supply voltage drops below the V_{LT} threshold.

*Table 4-5 Power Reset Limits**

DC VOLTAGE	THRESHOLD
+5V	V_{LT} = 4.25V minimum V_{HT} = 4.365V maximum
+12V	V_{LT} = 10.16V minimum V_{HT} = 10.48V maximum

* Includes a 75 mV peak-to-peak ripple on +5V or 150 mV peak-to-peak ripple on +12V to maximize or minimize values.

4.5.3 Power Requirements

This section lists the various power specifications for the Quantum Atlas 10K II Ultra 160/m SCSI hard disk drives.

4.5.3.1 Power Consumption

	Maximum Watts			
	9.2 GB	18.4 GB	36.7 GB	73.4 GB
Startup (Peak)	38.8 W*	38.8 W*	38.8 W*	47.6 W*
40% Seek (ECMA)	15.0 W	15.0 W	16.7 W	20.7 W
Idle	8.6 W	8.6 W	9.9 W	14.7 W

* Quantum Corporation's startup peak current measurements are based on a 2.5 ms sampling rate. The industry typical sampling rate is 20ms and, therefore, the Quantum stated startup peak power data appear to be higher than the industry typical numbers.

NOTE

ECMA Seek Rate is defined in ECMA standard 74. It corresponds to 52 random seeks per second.

4.5.3.2 DC Voltage Tolerances

Nominal Voltage	Voltage Tolerance	Absolute Limits
5 Vdc Supply	± 5%	4.75 - 5.25 Vdc
12 Vdc Supply	± 5% (±6% transient during spin-up)	11.4 - 12.6 Vdc

CAUTION

Polarity must not be reversed at any time. Reversing polarity at any time on either power line may damage components. Voltage is measured at the drive power connector.

4.5.3.3 Current Requirements

Typical currents are at nominal supply voltages.

Spin-Up Mode

	Typical / Maximum Current			
	9.2 GB	18.4 GB	36.7 GB	73.4 GB
5 Vdc Supply, Avg.	740.2 / 810* mA	686.0 / 820* mA	731.8 / 820* mA	719.5 / 820* mA
12 Vdc Supply, Avg.	2524.5 / 2850 mA, peak	2475.1 / 2860 mA, peak	2507.4 / 2860 mA, peak	3221.5 / 3740 mA, peak

* This value achieved during calibration.

Idle Mode

	Maximum			
	9.2 GB	18.4 GB	36.7 GB	73.4 GB
5 Vdc Supply, Avg.	570 mA	580 mA	580 mA	570 mA
12 Vdc Supply, Avg.	490 mA	610 mA	610 mA	1030 mA

ECMA Seek Mode

(ECMA Seeking)	Maximum			
	9.2 GB	18.4 GB	36.7 GB	73.4 GB
5 Vdc Supply, Avg.	600 mA	610 mA	610 mA	590 mA
12 Vdc Supply, Avg.	960 mA	1080 mA	1080 mA	1450 mA

NOTE

ECMA Seek Rate is defined in ECMA standard 74. It corresponds to 52 random seeks per second.

4.5.3.4 Current Transients

Spin-Up Mode

Max. Positive Rate of Change	+2.0 A in 200 μ s
Max. Negative Rate of Change	-2.0 A in 50 μ s
12 Vdc Supply Regulations	\pm 5%

Active Mode (ECMA Seeking)

Max. Positive Rate of Change	+1.0 A in 200 μ s
Max. Negative Rate of Change	-1.0 A in 200 μ s
12 Vdc Supply Regulations	\pm 5%

NOTE

ECMA Seek Rate is defined in ECMA standard 74. It corresponds to 52 random seeks per second.

NOTE

Instantaneous Peak Current transients of < 3A for < 50 μ s may occur during spin-up.

4.5.3.5 Supply Ripple Voltage and Noise

Supply Noise Voltage is the AC component from 100 KHz to 50 MHz measured with the supply operating the drive.

5 Vdc Supply Noise	150 mV p - p maximum
12 Vdc Supply Noise	250 mV p - p maximum

4.5.3.6 DC Voltage Monitoring

Unsafe Voltage Detection Level	
5 Vdc Supply Too Low	4.50 V nominal trip point
12 Vdc Supply Too Low	10.1 V nominal trip point

4.6 ACOUSTICS

Table 4-6 specifies the acoustical characteristics of the Quantum Atlas 10K II Ultra 160/m SCSI hard disk drive. Acoustic measurements are taken in an anechoic chamber with background noise < 25 dBA.

Table 4-6 Acoustical Characteristics

	Acoustic Emissions (Typ / Max)			
	9.2 GB	18.4 GB	36.7 GB	73.4 GB
Typical				
Active (ECMA Seeking)	4.4 / 4.7 dB	4.4 / 4.7 dB	4.4 / 4.7 dB	4.9 / 5.2 dB
Idle	3.9 / 4.2 dB	3.9 / 4.2 dB	3.9 / 4.2 dB	4.3 / 4.6 dB

NOTE

ECMA Seek Rate is defined in ECMA standard 74. It corresponds to 52 random seeks per second.

4.7 MECHANICAL

Dimensions	9.1 GB	18.4 GB	36.7 GB	73.4 GB
Height (in / mm) (3-sigma)	1.0 inch / 25.72 mm ± 0.24	1.0 inch / 25.72 mm ± 0.24	1.0 inch / 25.72 mm ± 0.24	1.64 inch / 41.6 mm ± 0.24
Width (in / mm)	4.0 inch / 101.6 mm ± 0.25	4.0 inch / 101.6 mm ± 0.25	4.0 inch / 101.6 mm ± 0.25	4.0 inch / 101.6 mm ± 0.25
Length (in / mm)	5.75 inch / 146.5 mm ± 0.25	5.75 inch / 146.5 mm ± 0.25	5.75 inch / 146.5 mm ± 0.25	5.75 inch / 146.5 mm ± 0.25
Weight (lbs / kg, nominal)	1.4 / 0.64	1.4 / 0.64	1.5 / 0.68	2.4 / 1.1

NOTE

All dimensions are exclusive of any optional faceplate (bezel).

4.8 ENVIRONMENTAL CONDITIONS

Table 4-7 summarizes the environmental specifications of the Quantum Atlas 10K II Ultra 160/m SCSI hard disk drive.

Table 4-7 Environmental Specifications

PARAMETER	OPERATING	NON-OPERATING
Temperature (Non-condensing)	5° to 55°C (41° to 131°F)	-40° to 70°C (-40° to 158°F)
Temperature Gradient (Non-condensing)	20°C (68°F) per hour, max.	30°C (86°F) per hour, max.
Humidity ¹ (non-condensing) Maximum Wet Bulb Temperature	5% to 85%, RH 37°C (98.6°F)	5% to 95%, RH 40°C (104°F)
Humidity Gradient	30% per hour	30% per hour
Altitude ²	-198 to +3049 m (-650 to 10,000 feet)	-198 to +12,195 m (-650 to 40,000 feet)

NOTES:

- ¹ No condensation.
- ² Altitude is relative to sea level.

4.9 SHOCK AND VIBRATION

The Quantum Atlas 10K II Ultra 160/m SCSI hard disk drive can withstand levels of shock and vibration applied to any of its three mutually perpendicular axes, or principal base axis, as specified in Table 4-8. A functioning drive can be subjected to specified operating levels of shock and vibration. When a disk drive has been subjected to specified non-operating levels of shock and vibration, with power to the disk drive off, there will be no change in performance at power on.

When packed in its single-pack shipping container, the disk drive can withstand a drop of 36 inches onto a concrete surface on any of its six faces, three edges, or one corner. The 12-pack shipping container can withstand a drop from 30 inches onto a concrete surface on any of its six faces, three edges, or one corner.

Table 4-8 Shock and Vibration Specifications

PARAMETER	OPERATING	NON-OPERATING
Shock	15 G peak linear half-sine, 11 ms duration (3 axes) (READ only) 63/10 G peak linear half-sine, 2 ms duration (3 axes) (READ/WRITE)	Maximum Linear Shock (unpackaged): 1": Maximum 70 G, 11 ms duration, half-sine; 130 G 1 ms duration, half-sine; 220 G 2 ms duration, half-sine 1.6": Maximum 70 G, 11 ms duration, half-sine; 110 G 1 ms duration, half-sine; 150 G 2 ms duration, half-sine Maximum Linear Shock (packaged): Drop: from 91.4 cm (36.0 inches) for packages up to 20 lbs.; 76.2 cm (30.0 inches) for packages between 21 and 40 lbs. Drop on 6 faces, 3 edges, 1 corner.
Vibration	Sine/Swept, G, Peak to Peak: <u>Frequency (Hz)</u> <u>1.0"</u> <u>1.6"</u> 5 – 300 2.0 2.0 300 – 500 0.5 0.5	Swept Sine (G, Peak to Peak, 5 to 500Hz): 2.0
Vibration, Random	G ² /Hz: <u>Frequency (Hz)</u> <u>1.0"</u> <u>1.6"</u> 10 – 300 0.008 0.004 300 – 500 0.0012 0.0006	G ² /Hz: <u>Frequency (Hz)</u> 10 – 300 0.05 300 – 500 0.01

4.10 RELIABILITY

Component Life: 5 years

Preventive Maintenance: Not required

Start/Stop: 40,000 cycles (minimum, @ 25°C [77°F])

Drive reliability is closely related to the temperatures to which the disk drive is exposed. The operational design temperature ambient is 30°C (86°F). Temperature ambient above 50°C (122°F) or below 5°C (41°F) may decrease drive reliability.

Chapter 5

SCSI DESCRIPTION

This chapter contains an overview of SCSI command and status processing and a detailed description of the commands supported by the disk drives. The SCSI command system enables the initiator to instruct the drive to perform specific functions.

NOTE

In this manual, unless otherwise stated, **numerical values are given in decimal**. Hexadecimal numbers, such as opcodes, are always given with an "h" following, as in 5Ah except when entire data tables are in hexadecimal.

5.1 OVERVIEW OF THE SCSI COMMAND DESCRIPTIONS

The disk drives support the SCSI-3 commands listed in Table 5-1. The command categories are *sequential*, *normal*, and *immediate*.

Immediate commands are processed when received by the drive. In most cases, these commands do not require drive resources, do not change the state of the drive, and bypass the command queue (unless the command is tagged).

Sequential commands execute in the order that they arrive (unless accompanied by a HEAD OF QUEUE tag) and execute to completion before a subsequent command is activated.

Normal commands are allowed to execute concurrently (with the restriction that the drive executes overlapping writes in the order received). Normal commands are usually I/O commands.

Table 5-1 Supported SCSI Commands

Command	Operation Code	Category
CHANGE DEFINITION	40h	Sequential
FORMAT UNIT	04h	Sequential
INQUIRY	12h	Immediate
LOG SELECT	4Ch	Sequential
LOG SENSE	4Dh	Sequential
MODE SELECT (6)	15h	Sequential
MODE SELECT (10)	55h	Sequential
MODE SENSE (6)	1Ah	Sequential
MODE SENSE (10)	5Ah	Sequential
PERSISTENT RESERVATION IN	5Eh	Sequential
PERSISTENT RESERVATION OUT	5Fh	Sequential
READ (6)	08h	Normal
READ (10)	28h	Normal
READ BUFFER	3Ch	Sequential
READ CAPACITY	25h	Immediate
READ DEFECT DATA (10)	37h	Sequential
READ DEFECT DATA (12)	B7h	Sequential
READ LONG	3Eh	Sequential
READ SKIP MASK	E8h	Normal
REASSIGN BLOCKS	07h	Sequential
RECEIVE DIAGNOSTIC RESULTS	1Ch	Sequential

(continued...)

Table 5-1 Supported SCSI Commands (continued)

Command	Operation Code	Category
RELEASE	17h	Sequential
RELEASE (10)	57h	Sequential
REPORT LUNS	A0h	Sequential
REPORT DEVICE IDENTIFIER	A3h	Sequential
REQUEST SENSE	03h	Immediate
RESERVE	16h	Sequential
RESERVE (10)	56h	Sequential
REZERO UNIT	01h	Sequential
SEEK (6)	0Bh	Sequential
SEEK (10)	2Bh	Sequential
SEND DIAGNOSTIC	1Dh	Sequential
SET DEVICE IDENTIFIER	A4h	Sequential
START STOP UNIT	1Bh	Sequential
SYNCHRONIZE CACHE	35h	Sequential
TEST UNIT READY	00h	Immediate
VERIFY	2Fh	Normal
WRITE (6)	0Ah	Normal
WRITE (10)	2Ah	Normal
WRITE AND VERIFY	2Eh	Normal
WRITE BUFFER	3Bh	Sequential
WRITE LONG	3Fh	Sequential
WRITE SAME	41h	Sequential
WRITE SKIP MASK	EAh	Normal
XDREAD	52h	Normal
XDWRITE	50h	Normal
XPWRITE	51h	Normal

- NOTES:**
1. Relative Addressing is not supported by the drive. Therefore, in all I/O commands, the RelAdr bit must be 0.
 2. RESERVE and RELEASE are supported, as are third-party reservations. Extent reservations are not supported.
 3. The RECEIVE DIAGNOSTIC RESULTS and SEND DIAGNOSTIC DATA commands implement vendor-unique pages to test the drive during the manufacturing process. It is recommended that initiators specify only the non-page format variants of these commands (PF=0), except for page 0x40.

5.2 Command Descriptor Block

An initiator communicates with the drive by sending a 6-, 10-, or 12-byte Command Descriptor Block (CDB) that contains the parameters for the specific command. The SCSI command's operation code is always the first byte in the CDB and a control field is the last byte. For some commands, the CDB is accompanied by a list of parameters sent during the data-out buffer transfer. Figure 5-1 shows the format of a typical 6-byte CDB.

Table 5-2 contains a description of the CDB fields.

Bit	7	6	5	4	3	2	1	0
Byte								
0	Operation Code							
1	Reserved			(MSB)	Logical Block Address (LBA)			
2 - 3	Logical Block Address (LBA)							(LSB)
4	Transfer Length, Parameter List Length, or Allocation Length							
5	Control							

NOTE: Reserved fields in CDBs and Parameters have a value of 0; Reserved fields in states and other parameters sent to an initiator are set to 0.

Figure 5-1 Typical (6-Byte) Command Descriptor Block — Data Format

Table 5-2 Command Descriptor Block — Field Descriptions

Field	Description																		
Operation Code	<p>The first byte of a SCSI CDB contains an operation code. The operation code of the CDB has a <i>Group Code</i> field (bits 7-5) and a <i>Command Code</i> field (bits 4-0). The 3-bit Group Code field provides for eight groups of command codes. The 5-bit Command Code field provides for 32 command codes in each group. A total, therefore, of 256 possible operation codes exist. Operation codes are defined in the SCSI command standards. The group code for CDBs specified therein correspond to the length of the command descriptor as below:</p> <table border="0"> <thead> <tr> <th><u>Group Code</u></th> <th><u>Meaning</u></th> </tr> </thead> <tbody> <tr> <td>0</td> <td>6-byte commands</td> </tr> <tr> <td>1</td> <td>10-byte commands</td> </tr> <tr> <td>2</td> <td>10-byte commands</td> </tr> <tr> <td>3</td> <td>Reserved</td> </tr> <tr> <td>4</td> <td>16-byte commands</td> </tr> <tr> <td>5</td> <td>12-byte commands</td> </tr> <tr> <td>6</td> <td>Vendor specific</td> </tr> <tr> <td>7</td> <td>Vendor specific</td> </tr> </tbody> </table> <p>The operation code specifies the command being requested. The list of supported SCSI commands and their operation codes are contained in Table 5-1.</p>	<u>Group Code</u>	<u>Meaning</u>	0	6-byte commands	1	10-byte commands	2	10-byte commands	3	Reserved	4	16-byte commands	5	12-byte commands	6	Vendor specific	7	Vendor specific
<u>Group Code</u>	<u>Meaning</u>																		
0	6-byte commands																		
1	10-byte commands																		
2	10-byte commands																		
3	Reserved																		
4	16-byte commands																		
5	12-byte commands																		
6	Vendor specific																		
7	Vendor specific																		

(continued)

Table 5-2 Command Descriptor Block — Field Descriptions (continued)

Field	Description
Logical Block Address	<p>The 6-byte READ, SEEK, and WRITE Command Descriptor Blocks contain a 21-bit Logical Block Address. The 10-, 12-, and 16- Command Descriptor Blocks contain a 32-bit Logical Block Address.</p> <p>Commands that require additional parameter data specify the length of the Logical Block Address that is needed. See the specific command descriptions for more detailed information.</p> <p>Relative Addressing indicates a technique used to determine the next Logical Block Address to be operated on. The drive does not support Relative Addressing, it defaults to a value of 0, which specifies that the Logical Block Address specifies the first logical block of a range of logical blocks to be operated on by the command.</p>
Transfer Length	<p>The transfer length field normally specifies the number of sectors to be transferred between the initiator and the drive. For several commands, the transfer length indicates the number of bytes (not sectors) to be sent. For these commands, this field may be identified by a different name</p> <p>Commands that use <i>one byte</i> for the transfer length value allow up to 256 sectors of data to be transferred by one command. A transfer length value of 0 indicates that 256 sectors are to be sent. Transfer length values of 1 through 255 indicate the number of sectors to be transferred.</p> <p>Commands that use <i>multiple bytes</i> for the transfer length value function differently. A transfer length value of 0 indicates that no data transfer is to occur. Transfer length values of 1 or greater indicate the number of sectors to be transferred.</p>
Parameter List Length	<p>The Parameter List Length is used to specify the number of bytes sent during the data-out buffer transfer. This field is typically used for parameters that are sent to a drive (for example, mode, diagnostic, and log parameters). A parameter list length of 0 indicates that no data is to be transferred.</p>
Allocation Length	<p>The Allocation Length field specifies the maximum number of bytes that the initiator has allocated for returned data. The Allocation Length is used to limit the amount of data returned to the initiator.</p> <p>An Allocation Length of 0 indicates that no data is to be transferred from the drive to the initiator. The drive terminates the data-in buffer transfer when the specified number of bytes have been transferred to the initiator or when all available data has been transferred, whichever is less.</p>
Control Field	<p>The Control Field is the last byte of every Command Descriptor Block; its format is shown in Figure 5-2 and described in Table 5-3.</p>

Bit	7	6	5	4	3	2	1	0
Byte	Vendor Specific		Reserved			NACA	Flag	Link

Figure 5-2 Command Descriptor Block Control Field — Data Format

Table 5-3 Command Descriptor Block Control Field — Field Descriptions

Field	Description
Vendor Specific Bits	These bits must be 0.
NACA	Normal Auto-Contingent Allegiance - This bit must be zero to indicate that SCSI-2 Contingent Allegiance rules apply.
<u>Link Bit</u>	<p>A Link bit set to one signals that the initiator requests continuation of a task (I/O Process) across two or more SCSI commands. If the Link bit is one and the flag bit is zero, and the command completes successfully, the drive will continue the task and return a status of INTERMEDIATE and a service response of Linked Command Complete.</p> <p>If the Link bit and the Flag bit of the Control word are both set to one, and the drive completes a command with a status of INTERMEDIATE, the drive shall return a service response of Linked Command Complete (with Flag). Refer to Section 5.6 for Linked Commands description.</p>
<u>Flag Bit</u>	The Flag bit is used in conjunction with the Link Bit to notify the initiator in an expedient manner that a command has been completed. A flag bit set to 1 is valid only when the Link Bit is set to 1.

5.3 Status/Error Reporting

SCSI message-level errors are communicated by messages that are defined specifically for that purpose. SCSI command-level errors are communicated by a status that is returned by the drive during the STATUS phase. This phase occurs at the end of each command, unless the command is terminated by one of the following events:

- ABORT TASK SET message
- ABORT TASK message
- TARGET RESET message
- CLEAR QUEUE message
- Unexpected disconnect

The status code is contained in bits 1 through 5 of the status byte. Bits 0, 6, and 7 are reserved. Table 5-4 describes the status codes returned by the drive.

Table 5-4 Status Codes

Status Code	Definition	Meaning
00h	GOOD	The drive successfully completed the command.
02h	CHECK CONDITION	An Auto Contingent Allegiance (ACA) condition occurred.
08h	BUSY	<p>The drive cannot service the command at the moment, and its Command Descriptor Block has been discarded. The initiator can retry the command at a later time. This status is returned when:</p> <ul style="list-style-type: none"> • A non-tagged command is received and the logical unit's command queue is full (all internal command buffers are in use). • A disconnect privilege was not granted in the IDENTIFY message of a queue-tagged I/O process (Parallel SCSI). • A disconnect privilege was not granted in the IDENTIFY message of a non-tagged I/O process and a command from another initiator is currently active (Parallel SCSI). • A command is received while an auto-contingent allegiance condition exists for another initiator.
10h	INTERMEDIATE	This status is returned for every command (except the last) in a series of linked commands that was successfully completed. However, if the command is terminated with other than GOOD status (such as CHECK CONDITION, RESERVATION CONFLICT, OR BUSY), the INTERMEDIATE status is not returned and the series of linked commands and the task is ended.
18h	RESERVATION CONFLICT	Another initiator has reserved the drive. (This status is never returned for INQUIRY or REQUEST SENSE commands.)
28h	TASK SET FULL	The drive cannot service the command at the moment, and its Command Descriptor Block has been discarded. (Returned for a tagged command when all of the drive's internal command buffers are in use, or when a host sends a tagged command while an Auto Contingent Allegiance condition is pending for that initiator)

5.4 Auto Contingent Allegiance Condition and Contingent Allegiance Condition

Auto Contingent Allegiance condition (ACA) is the SCSI-3 term for the condition of the drive immediately after it returns a CHECK CONDITION status code. While it is in this condition, the drive:

- Keeps the status information (sense data) that characterizes the nature of, and the reason for, the command's failure.
- Suspends execution of all commands that were queued behind the command that failed (includes commands from all initiators, tagged or untagged).
- Discards and returns BUSY status for both tagged and untagged drive access commands from other initiators.

The Auto Contingent Allegiance condition remains in effect until cleared by one of the following:

- TARGET RESET message received from any initiator
- ABORT TASK SET message received from the allied initiator
- Subsequent command received from the allied initiator

After receipt of an ABORT TASK SET message from the “allegiated” initiator, the drive clears any queued command sourced by that initiator and processes the remaining queued commands.

For Parallel SCSI, as its first action following receipt of the CHECK CONDITION status code, the initiator sends a REQUEST SENSE command to retrieve the sense data. Any other command from the initiator causes the drive to set the sense key to NO SENSE and the additional sense code to NO ADDITIONAL SENSE INFORMATION. These actions happen prior to resuming execution of queued commands, or, if there are no commands queued, before processing a new command.

Also for Parallel SCSI, after receipt of the REQUEST SENSE command or any other command from the “allegiated” initiator, the drive resumes processing of the queued commands in the normal manner if the QErr bit in the MODE SELECT Control Mode page is set to 0, its default value. If the value of QErr is 1, the drive aborts all queued commands and generates a Unit Attention condition for initiators with commands in the queue. A Unit Attention condition is not generated for the “allegiated” initiator if commands belonging to it are flushed from the queue.

5.5 Extended Contingent Allegiance Condition

Under SCSI-2 processing, the Extended Contingent Allegiance condition extends Contingent Allegiance handling so the initiator has explicit control over the point at which suspended (queued) commands are resumed. It allows the initiator to send a REQUEST SENSE command and then send the non-tagged commands to access the drive's medium.

This action is provided to fix damage left by the failed command before permitting the drive to continue with processing the queued commands.

The Extended Contingent Allegiance condition is enabled and disabled by the EECA bit in the MODE SELECT Control Mode page. The Extended Allegiance Condition differs from the Contingent Allegiance in that, under the Extended Contingent Allegiance condition:

- The drive sends an INITIATE RECOVERY message to the allegiance initiator prior to returning the failed command's COMMAND COMPLETE message. If the initiator rejects the INITIATE RECOVERY message, the drive drops back to contingent allegiance level of operation for the current error.
- An ABORT message does not terminate the extended contingent allegiance condition.
- Non-tagged drive access commands from the "allegianced" initiator are accepted by the drive and take priority over all queued commands. QUEUE FULL status is returned if the host sends a tagged command.
- The drive delays resuming or aborting the queued commands until a RELEASE RECOVERY message is received from the initiator.

The extended contingent allegiance condition remains in effect until cleared by one of the following:

- Hard reset
- BUS DEVICE RESET message from any initiator
- RELEASE RECOVERY message from the "allegianced" initiator

5.6 Linked Commands

An I/O Process (task) may contain multiple commands that are linked together. The initiator communicates this condition of linked (or unlinked) commands by setting (or clearing) the Link bit of the Command Descriptor Block's control word. A linked command is one in which the Link bit in the Command Descriptor Block is set.

After successful completion of a linked command the drive sends an INTERMEDIATE status, followed by a LINKED COMMAND COMPLETE message. If the Flag bit was set in the Command Descriptor Block, the drive sends an INTERMEDIATE status, followed by a LINKED COMMAND COMPLETE (WITH FLAG) message. The drive then switches the bus to the command phase in order to receive the next command in the linked chain.

All commands in a linked chain are addressed to the same nexus and are part of a single task (I/O process). The drive defers any commands that are not part of the linked chain until the chain is complete. The last command in the chain has the Link bit cleared.

NOTE: Relative addressing is not supported by the Atlas 10K II disk drive.

5.7 DATA Transfer Command Components

Many of the SCSI commands cause data to be transferred between the initiator and the drive. The content and characteristics of this data are command-dependent. Table 5-5 lists the information transmitted for all of the commands.

The “Length in CDB” column of Table 5-5 identifies the Command Descriptor Block field used by the drive to determine how much command-related data are to be transferred. The units (bytes or logical blocks) for the different Length fields are implied by the Length Field Name as follows:

<u>Field Name</u>	<u>Units Implied</u>
Allocation Length	Bytes of data the drive is allowed to send to the initiator
Parameter List Length	Bytes of data the initiator has available for the drive
Transfer Length	Logical data sectors the initiator wants transferred or verified
Byte Transfer Length	Bytes of data the initiator wants transferred

The DATA OUT column in Table 5-5 lists the information passed to the drive by the initiator as part of the command. The DATA IN column lists the information sent to the initiator by the drive.

Numbers in parentheses after an item indicate the item’s length in bytes. In some cases, additional length information is communicated during the DATA phase. For example, a FORMAT UNIT Defect List Header contains a Defect List Length field that contains the total length of the Defect Descriptors that follow the Defect List Header. Table 5-5 does not include these cases.

Table 5-5 DATA-Phase Command Contents

Command	Length in CDB	Data Out (To Drive)	Data In (To Initiator)
CHANGE DEFINITION	0	---	---
FORMAT UNIT	0	Defect List Header (4) Initialization Pattern (6-8) Defect Descriptors	---
INQUIRY	Allocation	---	Standard Inquiry or a Vital Product Data page
LOG SELECT	Parameter List (must be 0)	---	---
LOG SENSE	Allocation	---	Log Page
MODE SELECT	Parameter List	Mode Parameter Header (4) Block Descriptor (8) Page(s)	---
MODE SENSE	Allocation	---	Block Descriptor (8)
PERSIST. RES. IN			
PERSIST. RES. OUT			
READ (6) (10)	Transfer	---	Data
READ BUFFER	Allocation	---	READ BUFFER Header (4) Mode-zero Buffer (512) or Section of Drive's DRAM or READ BUFFER Desc. (4)
READ CAPACITY	Allocation	---	READ CAPACITY data (8)
READ DEFECT DATA	Allocation	---	<ul style="list-style-type: none"> • Defect List (Hdr) (4) • Defect Descriptors
READ LONG	Byte Transfer (Must be 560)	---	<ul style="list-style-type: none"> • Data (512) • LBA Tag (2) • EDC (2) • ECC (44)
READ SKIP MASK			
REASSIGN BLOCKS	0	Defect List Header(4)	Defect Descriptors
RECEIVE DIAGNOSTIC RESULTS	Allocation	---	Diagnostic Page • ECC (25)
RELEASE	0	---	---
REPORT DEVICE IDENTIFIER			
REQUEST SENSE	Allocation	---	Sense Data (18)
RESERVE	0 (Extent List Option not supported)	---	---
REZERO UNIT	0	---	---

(continued)

Table 5-5 DATA-Phase Command Contents (continued)

Command	Length in CDB	Data Out (To Drive)	Data In (To Initiator)
SEEK (6)(10)	0	---	---
SEND DIAGNOSTIC	Parameter List	Diagnostic Page	---
SET DEVICE IDENTIFIER			
START STOP UNIT	0	---	---
SYNCHRONIZE CACHE	0	---	---
TEST UNIT READY	0	---	---
VERIFY	Transfer	Data	---
WRITE (6)(10)	Transfer	Data	---
WRITE AND VERIFY	Transfer	Data	---
WRITE BUFFER	Parameter List	Zeros (4) Mode-zero buffer (512) or Data to put into DRAM or Microcode image (262,144) or Microcode image (successive 16 KB pieces)	---
WRITE LONG	Byte Transfer (must be 560)	<ul style="list-style-type: none"> • Data (512) • LBA Tag (2) • EDC (2) • ECC (44) 	---
WRITE SAME	0	Data (1 logical sector)	---
WRITE SKIP MASK			
XDREAD			
XDWRITE			
XPWRITE			

5.8 SCSI COMMAND DESCRIPTIONS

The SCSI command descriptions that follow this page contain detailed information about the SCSI commands that are supported by the drive. Each description provides a Data Format and Field Descriptions for the Command Descriptor Block for the described command.

The commands are presented in alphabetic order, and each command starts on a new, odd-numbered page.

Common Fields

Several fields that are common to many commands are described here, rather than being repeated throughout the descriptions. These fields include:

Reserved – Reserved bits, fields, bytes, and code values are set aside for future standardization and must be set to 0. If the drive receives a command that contains non-0 bits in a reserved field or a reserved code value, the command is terminated with a CHECK CONDITION status and the sense key set to ILLEGAL REQUEST. However, there are some fields that are not checked for compatibility with older SCSI initiators.

Control – The Link Bit and Flag Bit are supported.

RelAdr – Not supported; **must be 0**.

5.9 CHANGE DEFINITION Command (40h)

The CHANGE DEFINITION command sets the SCSI compliance for disk drives with parallel SCSI interfaces to one of four different levels: SCSI-1, SCSI-1/CCS, SCSI-2 or SCSI-3.

Bit	7	6	5	4	3	2	1	0
Byte								
0	Operation Code (40h)							
1	Reserved							
2	Reserved							Save
3	Rsv'd	New Operating Definition (0, 1, 2, 3, or 4 are legal values for Parallel SCSI Disk Drives)						
4 - 7	Reserved							
8	Parameter Data Length							
9	Control							

Figure 5-3 CHANGE DEFINITION Command Descriptor Block — Data Format

Table 5-6 CHANGE DEFINITION — Field Descriptions

Field	Description
Save	Save Parameters. When this field is set to 1, the new operating definition is saved in the drive's non-volatile memory.
New Operating Definition	The values that can be supplied in the field are listed below: 0 Current Definition 1 SCSI-1 2 SCSI-1/CCS (Common Command Set) 3 SCSI-2 4 SCSI-3
Parameter Data Length	Length, in bytes, of the parameter data. Must be 0.

5.10 FORMAT UNIT Command (04h)

The FORMAT UNIT command formats the disk's storage media into initiator-addressable logical blocks according to initiator-defined options. This command ensures that the disk storage media is formatted so that all data sectors are accessible. Any data residing on the disk before this command is invoked is lost.

This command repairs damage left by previous WRITE LONG commands. The WRITE LONG command allows the initiator to deliberately corrupt a sector's ECC or EDC.

The FORMAT UNIT command updates the defective sector list, referred to as the Grown Defect List, that is maintained by the drive. As described below, the drive gives the initiator considerable control over this list. If certification is enabled, all initiator-addressable logical blocks are verified, and bad sectors are reassigned and added to the Grown Defect List.

NOTE

During FORMAT UNIT command processing, the drive ignores the Read/Write AWRE and ARRE bits (from Error Recovery Mode page) and does auto-revectoring as needed.

The drive is physically formatted when it leaves the manufacturing facility. Therefore, it is not necessary to reformat the drive before using it. If the logical sector size of the drive is changed using the Block Descriptor of the MODE SELECT Command, it is recommended (but not necessary) to reformat the drive.

The Immed bit that can be included in the FORMAT UNIT command parameter list allows the initiator to control whether the drive returns completion status either after fetching and validating the Command Descriptor Block and parameter list or after the FORMAT UNIT command completes.

The time required for the FORMAT UNIT command to format the disk primarily depends on the capacity of the drive. The number of defects detected and the number of defects already in the Primary and Grown Defect Lists also influence the time required to format the drive. If the Disable Certification bit is set to 1, formatting time is reduced.

Defective sectors on the drive are managed through two lists: the Primary Defect List and the Grown Defect List. The Primary Defect List is created when the drive is manufactured and is the drive's initial defect list. The Primary Defect List is not affected by the FORMAT UNIT command. Sectors listed in it are revectoring by the drive. The Grown Defect List contains a list of the sectors that have gone bad since the drive's primary list was generated.

Defects are communicated in a data structure referred to as a Defect Descriptor. The SCSI specification defines several Defect Descriptor formats. The FORMAT UNIT command recognizes defect descriptors in Block format, Bytes From Index format, and Physical Sector format. (There is one exception: a value of FFFFFFFFh in either the defect Bytes From Index field or the Defective Sector Number fields is ignored).

FORMAT UNIT Command (04h) (continued)

Bit Byte	7	6	5	4	3	2	1	0
0	Operation Code (04h)							
1	Reserved			Fmt Data	Cmp List	Defect List Format		
2	Vendor-Specific							
3 - 4	Interleave							
5	Control							

Figure 5-4 FORMAT UNIT Command Descriptor Block — Data Format

Table 5-7 FORMAT UNIT Command — Field Descriptions

Field	Description
FmtData	Format Data. When the Format Data value is 1, it indicates that a data-out buffer transfer occurs as part of the command execution. The FORMAT UNIT Parameter List (consisting of a Defect List Header and, optionally, a number of Block Format Defect Descriptors) is passed to the device during this phase. When the Format Data value is 0, the data-out buffer transfer does not occur.
CmpLst	Complete List. When the Complete List value is 1, the drive deletes its current Grown Defect List and starts a new one, containing the Logical Block Numbers listed in the defect list supplied by the initiator with this command. Defective Logical Block Numbers identified during this format operation are added to the list, creating a new Grown Defect List. When the Complete List value is 0, the drive adds initiator-supplied and newly found defective Logical Block Numbers to the existing Grown Defective List.
Defect List Format	The Defect List Format value specifies the defect descriptor passed by the initiator to the drive when the Format Data value is 1. Acceptable values are: 000b (Block [or Sector] Format) 100b (Index Format) 101b (Physical Sector Format).
Vendor-Specific	Not supported. Must be 0.
Interleave	Not supported. Ignored by the drive.

FORMAT UNIT Command (04h) (continued)

5.10.1 Five Forms of FORMAT UNIT Commands

Five different forms of the FORMAT UNIT command are supported through different combinations of the Format Data bit, Complete List bit, and the information in the Defect List Length field. (Refer to *FORMAT UNIT Defect Header List* for a description of the Defect List Length field.) The different command forms give the initiator control over the contents of the Grown Defect List. Table 5-8 describes the contents of the Grown Defect List after FORMAT UNIT command execution.

Table 5-8 *FORMAT UNIT Command Supported Options*

FmtData	CmpLst	Defect List Length	Contents of Grown Defect List after FORMAT UNIT Command Execution
0	0	N/A	All sectors found to be bad, but not listed in the Primary Defect List or Grown Defect List, are added to the Grown Defect List.
1	0	0	Same as above.
1	1	0	The existing Grown Defect List is discarded. A new Grown Defect List is generated, containing all sectors found to be bad but not listed in the Primary Defect List.
1	0	>0	The Grown Defect List contains: <ul style="list-style-type: none"> • Original Grown Defect List. • All sectors found to be bad, but not listed in the Primary Defect List. • A list of sectors supplied by the initiator. These sectors are passed in defect descriptors in the data-out buffer transfer that occurs as part of the FORMAT UNIT command.
1	1	>0	Same as the case above, except that the current Grown Defect List is discarded before formatting begins.

5.10.2 FORMAT UNIT Parameter List

The FORMAT UNIT Parameter List (Figure 5-5) consists of a Defect List Header (Figure 5-6) followed by one or more Defect Descriptors. Descriptors are either four bytes or eight bytes in length.

Bit	7	6	5	4	3	2	1	0
Byte								
0 - 3	Defect List Header							
	Initialization Pattern Descriptor (if any)							
	Defect Descriptors (if any)							

Figure 5-5 *FORMAT UNIT Parameter List — Data Format*

*FORMAT UNIT Command (04h) (continued)***5.10.2.1 FORMAT UNIT Defect List Header**

The FORMAT UNIT Defect List Header (Figure 5-6) provides several optional format control bits to give the initiator more control over the defect lists. Table 5-9 provides descriptions of the data fields in the Defect List Header.

Bit	7	6	5	4	3	2	1	0
Byte								
0	Reserved							
1	FOV	DPRY	DCRT	STPF	IP	DSP	Immed	VS
2 - 3	Defect List Length							

Figure 5-6 FORMAT UNIT Defect List Header — Data Format

Table 5-9 FORMAT UNIT Defect List Header — Field Descriptions

Name	Description
FOV	The Format Options Valid bit indicates that the remaining option bits in this byte are valid. If this bit is not set, the remaining bits, except Immed, are ignored.
DPRY	The Disable Primary bit specifies whether the drive should ignore the Primary Defect List during the format to identify defective areas on the media. The drive's default is 0, indicating that the drive should replace sectors found in the Primary Defect List during the format.
DCRT	The Disable Certification bit specifies whether the drive should certify the media during the format. The drive's default is 0, indicating that the drive should certify the media.
STPF	The Stop Format bit specifies the error conditions under which the drive is to terminate the format. The state of the bit is ignored and the drive acts as if this bit is set. If either the Grown Defect List or the Primary Defect List is not found, the format operation terminates with a CHECK CONDITION status and a sense key of MEDIUM ERROR.
IP	The Initialization Pattern bit signals the drive that the FORMAT UNIT Parameter List contains initialization pattern information. An Initialization Pattern bit of 1 indicates that an initialization pattern descriptor (Figure 5-7) is included in the FORMAT UNIT parameter list immediately following the Defect List Header. An Initialization Pattern bit of 0 indicates that an initialization pattern descriptor (Figure 5-7) is not included in the FORMAT UNIT parameter list, and the drive will use its default initialization pattern of all zeros.

(continued)

*FORMAT UNIT Command (04h) (continued)**Table 5-9 FORMAT UNIT Defect List Header — Field Descriptions (continued)*

Name	Description
DSP	The Disable Saving Parameter prohibits the drive from preserving MODE SELECT parameters received while the Format is in progress. This bit is ignored, and the drive acts as if it were set.
Immed	<p>When the Immediate bit is 0, the drive returns a status after the format operation completes. If the disconnect privilege is granted in the accompanying IDENTIFY message, the drive disconnects from the bus after fetching and validating the Command Descriptor Block and the FORMAT UNIT Parameter List. If the disconnect privilege is not granted, the drive stays connected to the bus during the entire FORMAT UNIT command.</p> <p>When the Immediate bit is set to 1, the drive returns a status after the Command Descriptor Block and FORMAT UNIT Parameter List have been fetched and validated. In this case, the drive ignores the IDENTIFY message's disconnect privilege bit and remains connected to the bus while it validates the Command Descriptor Block and Parameter List. It returns the completion status before it disconnects. The drive then proceeds with the format.</p>
VS	Vendor-specific. Must be 0.
Defect List Length	Gives the length in bytes of the Defect Descriptors that follow. A value of 0 is valid and means that no Defect Descriptors follow. A CHECK CONDITION status is returned if this value is not a multiple of the defect descriptor size. The Defect List Length is equal to four times the number of defect descriptors if Block format is used, or eight times if Bytes From Index format or Physical Sector format is used.

5.10.2.2 FORMAT UNIT Defect Descriptor — Block Format

Each descriptor specifies a 4-byte defective sector address of the sector that contains the defect.

Bit	7	6	5	4	3	2	1	0
Byte	Defective Block Address							
0 - 3								

Figure 5-7 FORMAT UNIT Defect Descriptor — Block Format

*FORMAT UNIT Command (04h) (continued)***5.10.2.3 FORMAT UNIT Defect Descriptor — Physical Sector and Bytes From Index Format**

The Physical Sector defect descriptor specifies the location of a defect that is the length of a sector. The Bytes From Index defect descriptor specifies the location of a defect that is no more than eight bytes long.

Each descriptor is comprised of the bytes shown in Figure 5-8. The Cylinder Number of Defect is the most significant part of the address, and the Defective Sector Number or Defect Bytes From Index is the least significant part of the address. A value of FFFFFFFFh in bytes 4 to 7 is ignored.

Bit Byte	7	6	5	4	3	2	1	0
0 - 2	Cylinder Number of Defect							
3	Head Number of Defect							
4 - 7	Defective Sector Number or Defect Bytes from Index							

Figure 5-8 FORMAT UNIT Defect Descriptor — Physical Sector and Bytes From Index Format

5.10.2.4 FORMAT UNIT Initialization Pattern Descriptor

The Initialization Pattern option specifies that the logical blocks on the drive medium shall contain a specified initialization pattern. The Initialization Pattern descriptor is sent to the drive as part of the FORMAT UNIT parameter list. The contents of the descriptor fields are described in Table 5-10.

Bit Byte	7	6	5	4	3	2	1	0
0	IP Modifier = 0		Reserved					
1	IP Type							
2 - 3	IP Length							
4 - n	IP							

Figure 5-9 FORMAT UNIT Initialization Pattern Descriptor — Data Format

*FORMAT UNIT Command (04h) (continued)**Table 5-10 FORMAT UNIT Initialization Pattern Descriptor — Field Descriptions*

Name	Description
IP Modifier	The Initialization Pattern Modifier must be 0.
IP Type	The Initialization Pattern Type field (Table 5-11) indicates the type of pattern the drive uses to initialize each logical sector within the initiator-accessible portion of the medium. All bytes within a logical sector are written with the initialization pattern.
IP Length	The Initialization Pattern Length field indicates the number of bytes contained in the Initialization Pattern. The valid lengths (when IP Type = 1) are two or four bytes.
IP	Initialization Pattern.

Table 5-11 FORMAT UNIT Initialization Pattern Type

Initialization Pattern Type	Description
00h	Use default pattern. ¹
01h	Repeat the initialization pattern as required to fill the logical sector. ²
02h – FFh ³	Reserved.

- NOTES:**
- 1 If the Initialization Pattern Length is not 0, 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.
 - 2 If the Initialization Pattern Length is 0, 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.
 - 3 If the Initialization Pattern Type is not 0 or 1, the drive terminates the command with a CHECK CONDITION status. The Sense Key is set to ILLEGAL REQUEST, and the Additional Sense Code is set to INVALID FIELD IN PARAMETER LIST.

5.11 INQUIRY Command (12h)

The INQUIRY command allows the initiator to determine the kind of SCSI devices attached to its SCSI bus. It causes a device that is attached to a SCSI bus to return information about itself. The drive identifies itself as a Direct Access Storage Device that implements the applicable interfacing protocol. The drive does not need to access its storage medium to respond to the inquiry.

The drive can provide two categories of data in response to an INQUIRY command: Standard Inquiry Data and Vital Product Data. Standard Inquiry Data contains basic data about the drive, and Vital Product Data comprises several pages of additional data. Each Vital Product Data page requires a separate INQUIRY command from the initiator.

An INQUIRY command is not affected by, nor does it clear, a Unit Attention condition.

Bit	7	6	5	4	3	2	1	0
Byte								
0	Operation Code (12h)							
1	Reserved						CmdDt	EVPD
2	Page Code or Operation Code							
3	Reserved							
4	Allocation Length							
5	Control							

Figure 5-10 INQUIRY Command Descriptor Block — Data Format

Table 5-12 INQUIRY Command Descriptor Block — Field Descriptions

Field	Description
CmdDt	Command Support Data. If CmdDt = 0 and EVPD (see below) = 0, the drive returns the Standard Inquiry Data. If CmdDt = 1 with EVPD = 0, the drive returns the Command Data specified by Page Code/Operation.
EVPD	Enable Vital Product Data. If EVPD = 0 and CmdDt (see above) = 0, the drive returns the Standard Inquiry Data. If EVPD = 1 and CmdDt = 0, the drive returns the Vital Product Data Page specified by Page Code/Operation Code.

(continued)

INQUIRY Command (12h) (continued)

Table 5-12 INQUIRY Command Descriptor Block — Field Descriptions (continued)

Field	Description
Page Code or Operation Code	Specifies the Vital Product Data Page or Command Support Data which is to be returned by the drive when EVPD is set. Specifies the SCSI Operation Code for command support data to be returned by the drive when CmdDt is set. A CHECK CONDITION status is returned if this field specifies an unsupported Page or Operation Code or if both EVPD and CmdDt are set. Table 5-14 lists the Page Codes for the Vital Product Pages supported by the drive. Table 5-18 lists the operation codes for the CmdDt information supported by the drive.
Allocation Length	Specifies the number of bytes of inquiry information the drive is allowed to return to the initiator during the command's data-in buffer transfer. Error status is not returned if the value in this field truncates the requested information.

*INQUIRY Command (12h) (continued)***5.11.1 Standard Inquiry Data Page**

The Standard Inquiry Data Page is returned by the drive in response to the INQUIRY command if EVPD = 0.

Bit Byte	7	6	5	4	3	2	1	0
0	Peripheral Qualifier = 0			Peripheral Device Type = 0 (Direct Access Device)				
1	RMB = 0	Device Type Modifier = 0						
2	ISO Version = 0		ECMA Version = 0			ANSI Version = 3 (SCSI-3)		
3	AERC = 0	TrmTsk = 0	Norm ACA = 0	Rsv'd	Response Data Format = 2			
4	Additional Length = 5Bh							
5	Reserved							
6	Rsv'd	Rsv'd	Port = 0	MultiP = 0	Mchngr = 0	ACKQ- REQQ= 0	Addr32	Addr16
7	RelAdr = 0	Wbus32 = 0	Wbus16	Sync = 1	Linked = 1	TransDis = 1	CmdQue = 1	SftRe = 0
8 - 15	Vendor Identification "QUANTUM"							
16 - 31	Product Identification							
32 - 35	Product Revision Level							
36 - 47	Drive Serial Number							
48 - 51	Reserved							
52 - 53	SCSI Hardware Revision Number							
54	Disk Controller Hardware Revision Number							
55	Electronics Pass Number							
56	Reserved			Clocking		QAS	IUS	
57 - 95	Reserved							

Figure 5-11 Standard Inquiry Data Page — Data Format

INQUIRY Command (12h) (continued)

Table 5-13 Standard Inquiry Data Page — Field Descriptions

Field Name	Value	Description
Peripheral Qualifier	0	Non-zero if initiator selects an invalid logical unit.
Peripheral Device Type	0	0 indicates that this is a Direct Access Device.
ANSI Version	3	ANSI SCSI Level 3 (SCSI-3) is supported.
AERC	0	Asynchronous Event Reporting is not supported.
NormACA	0	Does not support setting NACA in CDB Control word.
Port	0	Only used when MultiP = 1.
MultiP	0	This field set to 1 if it is a multiport device.
Mchngr	0	Not embedded in or attached to a medium changer
ACKREQQ	0	Device supports REQ and ACK data handshake on Q cable.
Addr32	0	32-bit wide SCSI address
Addr16	0/1	16-bit wide SCSI address
TrmTSK	0	TERMINATE TASK Function is not supported
Response Data Format	2	This Standard Inquiry Data is in the format specified in the International Standard.
RelAdr	0	Relative Addressing is not supported.
WBus 32	0	32-bit wide bus is not supported.
WBus 16	0 or 1	The WBus bit is 1 if the drive supports 16-bit data transfer.
The		bit is 0 if the drive only supports 8-bit transfer.
Sync	1	The drive supports Synchronous Data Transfers.
Linked	1	Linked Commands are supported.
TransDis	1	Indicates that the drive supports CONTINUE TASK and TARGET TRANSFER DISABLE messages.
CmdQue	1	The drive supports Tagged Command Queuing.
SftRe	0	The drive implements the hard reset option in response to assertion of the SCSI Bus reset line.
VS	0	Vendor Specific
Product Identification		The value in this field is: ATLAS10KII-9WLS, ATLAS10KII-18 WLS, ATLAS10KII-36WLS ATLAS10KII-72WLS if the <i>parallel</i> drive supports 16-bit data transfers; or ATLAS10KII-9SCA, ATLAS10KII-18SCA, ATLAS10KII-36SCA ATLAS10KII-72SCA if the <i>parallel</i> drive supports 16-bit data transfers through SCA. Note: Product Identification may vary from the above, depending on OEM customer specifications.

NOTE: Vendor Information, Product Identification, and Product Revision Level are returned as shown in Figure 5-10.

INQUIRY Command (12h) (continued)

Table 5-13 Standard Inquiry Data Page — Field Descriptions (continued)

Field Name	Value	Description
Clocking	11	The contents of this field define the setting of the double-edge clocking option (See Chapter 6). Note that the clocking field does not apply to asynchronous transfers of data. Options available are:
		Code Description
		00b Indicates the device server supports only Single Transition (ST)
		01b Indicates the device server supports only Double Transition (DT)
		10b Reserved
		11b Indicates the device server supports ST and DT
QAS	0	Quick Arbitrate Support. If the value of this field is 1, it indicates that the device server supports the quick arbitrate feature. A value of 0 indicates that the device server does not support this feature. 0 is default setting.
IUS	0	Information Unit Supported. If the value of this field is 1, it indicates that the device server supports information units. A value of 0 indicates that the device server does not support information units. Default is 0.

5.11.2 Vital Product Data Pages

The Vital Product Data pages that can be returned by the drive are described in the following paragraphs in the sequence shown in Table 5-14.

5.11.2.1 Supported Vital Product Data Pages Page (00h)

The Supported Vital Product Data Pages page (Figure 5-12) provides a directory of the Vital Product Data Pages that are supported by the drive. Table 5-14 lists the supported pages.

Bit	7	6	5	4	3	2	1	0
Byte								
0	Peripheral Qualifier = 0			Peripheral Device Type = 0 (Direct Access Device)				
1	Page Code (00h)							
2	Reserved							
3	Page Length = 08h							
4 - 10	Supported Page List (see Table 5-14)							

Figure 5-12 Supported Vital Product Data Pages Page — Data Format

Table 5-14 Vital Product Data — Page Codes

Page Code	Description	Size In Bytes
00h	Supported Vital Product Pages Page	11
80h	Unit Serial Number Page	16
81h	Implemented Operating Definition Page	9
82h	ASCII Implemented Operating Page	32
83h	Device Identification Page	28
C0h	Firmware Revision Page	40
C1h	Quantum Manufacturing Number Page	20
C4h	Negotiated Rate Information Page	8

*INQUIRY Command (12h) (continued)***5.11.2.2 Unit Serial Number Page (80h)**

The Unit Serial Number page contains the drive's PCB Serial Number and the HDA Serial Number.

Bit	7	6	5	4	3	2	1	0
Byte								
0	Peripheral Qualifier = 000b			Peripheral Device Type = 000h (Direct Access Device)				
1	Page Code (80h)							
2	Reserved							
3	Page Length = 0Ch							
4 - 15	HDA Serial Number							

Figure 5-13 Unit Serial Number Page — Data Format

Table 5-15 Unit Serial Number Page — Field Descriptions

Field	Definition
HDA Serial Number	A 12-character ASCII representation of the drive's HDA serial number

*INQUIRY Command (12h) (continued)***5.11.2.3 Implemented Operating Definition Page (81h)**

The Implemented Operating Definition page reflects the current and available operation definitions.

Bit	7	6	5	4	3	2	1	0
Byte								
0	Peripheral Qualifier = 0			Peripheral Device Type = 0 (Direct Access Device)				
1	Page Code (81h)							
2	Reserved							
3	Page Length = 05h							
4	Rsv'd	Current Definition						
5	SavImp = 1	Default Definition is SCSI-3 (04h)						
6	SavImp = 1	Default Definition is SCSI-3 (04h)						
7 *	SavImp = 1	Default Definition is SCSI-2 (03h)						
8	SavImp = 1	Default Definition is SCSI-1CCS (02h)						

Figure 5-14 Implemented Operating Definition Page — Data Format

Table 5-16 Implemented Operating Definition Page — Field Descriptions

Field	Definition
SavImp	Save Implemented. If = 0, this bit indicates that the corresponding operation definition cannot be saved. When = 1, this bet indicates that the corresponding operating definition can be saved.

INQUIRY Command (12h) (continued)

5.11.2.4 ASCII Implemented Operating Definition Page (82h)

The ASCII Implemented Operating Definition page returns the character string's length (1Bh) in byte 4, followed by the appropriate character string ("SCSI-3, SCSI-2, SCSI-1/CCS" in bytes 5 through 31 for Parallel SCSI.

Bit Byte	7	6	5	4	3	2	1	0
0	Peripheral Qualifier = 0			Peripheral Device Type = 0 (Direct Access Device)				
1	Page Code (82h)							
2	Reserved							
3	Page Length = 1Ch							
4	ASCII String Length = 1Bh							
5 - 31	"SCSI-3, SCSI-2, SCSI-1/CCS" + byte of 0							

Figure 5-15 ASCII Implemented Operating Definition Page — Data Format

*INQUIRY Command (12h) (continued)***5.11.2.5 Device Identification Page (83h) (SCSI-3 ONLY)**

The Device Identification Page provides the means to retrieve zero or more identification descriptors that apply to the logical unit.

Bit	7	6	5	4	3	2	1	0
Byte								
0	Peripheral Qualifier (0)			Peripheral Device Type (0) (Direct Access Device)				
1	Page Code (83h)							
2	Reserved							
3	Page Length (24h for SCSI)							
4	Reserved				Code Set (1)			
5	Reserved		Association (0)		Identifier Type (2)			
6	Reserved							
7	Identifier Length (8h)							
8 – 15	EUI-64 (Value Stored in Configuration Page 30h)							
16	Reserved				Code Set (2)			
17	Reserved		Association (0)		Identifier Type (1)			
18	Reserved							
19	Identifier Length (14h)							
20 – 27	Vendor Identifier ("Quantum ")							
28 – 39	Drive Serial Number							

Figure 5-16 Device Identification Page — Data Format

INQUIRY Command (12h) (continued)

Table 5-17 Device Identification Page — Field Descriptions

Field	Definition														
Code Set	<p>This field specifies the code set used for the Identifier field. Applicable values are:</p> <table border="1"> <thead> <tr> <th>Value</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0h</td> <td>Reserved</td> </tr> <tr> <td>1h</td> <td>The Identifier field contains binary values</td> </tr> <tr> <td>2h</td> <td>The Identifier field contains ASCII graphic codes (code values 20h through 7Eh)</td> </tr> <tr> <td>3h – Fh</td> <td>Reserved</td> </tr> </tbody> </table>	Value	Description	0h	Reserved	1h	The Identifier field contains binary values	2h	The Identifier field contains ASCII graphic codes (code values 20h through 7Eh)	3h – Fh	Reserved				
Value	Description														
0h	Reserved														
1h	The Identifier field contains binary values														
2h	The Identifier field contains ASCII graphic codes (code values 20h through 7Eh)														
3h – Fh	Reserved														
Association	<p>This field specifies the entity with which the Identifier field is associated. Applicable values are:</p> <table border="1"> <thead> <tr> <th>Value</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0h</td> <td>The Identifier field is associated with the addressed physical or logical device.</td> </tr> <tr> <td>1h</td> <td>The identifier field is associated with the port that received the request.</td> </tr> <tr> <td>2h – 3h</td> <td>Reserved</td> </tr> </tbody> </table>	Value	Description	0h	The Identifier field is associated with the addressed physical or logical device.	1h	The identifier field is associated with the port that received the request.	2h – 3h	Reserved						
Value	Description														
0h	The Identifier field is associated with the addressed physical or logical device.														
1h	The identifier field is associated with the port that received the request.														
2h – 3h	Reserved														
Identifier Type	<p>This field specifies the format and assignment authority for the identifier. Values in this field are:</p> <table border="1"> <thead> <tr> <th>Value</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0h</td> <td>No assignment authority was used; there is no guarantee that the identifier is globally unique (vendor-specific).</td> </tr> <tr> <td>1h</td> <td>The first 8 bytes of the Identifier field represent the Vendor ID.</td> </tr> <tr> <td>2h</td> <td>The Identifier field contains an IEEE Extended Unique Identifier, 64-bit (EUI-64). The Identifier Length field (Byte 7) is set to 8.</td> </tr> <tr> <td>3h</td> <td>Not applicable; for Fibre Channel devices.</td> </tr> <tr> <td>4h</td> <td>Not applicable; if the Association value = 1h, the value of the Identifier contains a 4-byte, binary number that identifies the port relative to the other ports in the device.</td> </tr> <tr> <td>5h - Fh</td> <td>Reserved</td> </tr> </tbody> </table>	Value	Description	0h	No assignment authority was used; there is no guarantee that the identifier is globally unique (vendor-specific).	1h	The first 8 bytes of the Identifier field represent the Vendor ID.	2h	The Identifier field contains an IEEE Extended Unique Identifier, 64-bit (EUI-64). The Identifier Length field (Byte 7) is set to 8.	3h	Not applicable; for Fibre Channel devices.	4h	Not applicable; if the Association value = 1h, the value of the Identifier contains a 4-byte, binary number that identifies the port relative to the other ports in the device.	5h - Fh	Reserved
Value	Description														
0h	No assignment authority was used; there is no guarantee that the identifier is globally unique (vendor-specific).														
1h	The first 8 bytes of the Identifier field represent the Vendor ID.														
2h	The Identifier field contains an IEEE Extended Unique Identifier, 64-bit (EUI-64). The Identifier Length field (Byte 7) is set to 8.														
3h	Not applicable; for Fibre Channel devices.														
4h	Not applicable; if the Association value = 1h, the value of the Identifier contains a 4-byte, binary number that identifies the port relative to the other ports in the device.														
5h - Fh	Reserved														

*INQUIRY Command (12h) (continued)***5.11.2.6 Firmware Revision Page (C0h)**

The Firmware Revision page reflects the current firmware revision and date and time when it was built.

Bit	7	6	5	4	3	2	1	0
Byte								
0	Peripheral Qualifier = 0			Peripheral Device Type = 0 (Direct Access Device)				
1	Page Code (C0h)							
2	Reserved							
3	Page Length = 24h							
4 - 19	Full Firmware Version Blank Fill (e.g., "T382 ")							
20 - 39	Build Date and Time Blank Fill (e.g., "Jan 19 1999 07:18:15")							

Figure 5-17 Firmware Revision Page — Data Format

INQUIRY Command (12h) (continued)

5.11.2.7 Quantum Manufacturing Information Page (C1h)

The Quantum Manufacturing Information Page reflects the part number assigned by Quantum’s manufacturing group in accordance with vendor requirements.

Bit	7	6	5	4	3	2	1	0
Byte								
0	Reserved							
1	Page Code (C1h)							
2	Reserved							
3	Page Length = 10h							
4 - 19	Spaces or Quantum-Unique Information							

NOTE: The value in the Bytes 4 - 19 field is assigned by Quantum Corporation

Figure 5-18 Quantum Manufacturing Information Page — Data Format

*INQUIRY Command (12h) (continued)***5.11.2.8 Negotiated Rate Information Page (C4h)**

Bit Byte	7	6	5	4	3	2	1	0
0	Peripheral Qualifier = 0			Peripheral Device Type = 0 (Direct Access Device)				
1	Page Code (C4h)							
2	Reserved							
3	Page Length = 4h							
4	Transfer Period Factor							
5	REQ / ACK Offset							
6	Rsv'd	AIP	AIPD	AIPCE	TWE	PARL	FAIR	SYNC
7	Reserved					QAS	DT	IU

Figure 5-18a Negotiated Rate Information Page — Data Format

INQUIRY Command (12h) (continued)

Table 5-17A Negotiated Rate Information Page — Field Descriptions

Field	Definition
Transfer Period Factor	Applicable values of this field are provided in Appendix C.
REQ / ACK Offset	This is the negotiated value of Req / Ack offset between the initiator and the target. For Single Transition (ST) data transfers, the REQ/ACK Offset is the maximum number of REQ assertions allowed to be outstanding before a corresponding ACK assertion is received at the target: Data transfer size may be 1 or 2 bytes. For Double Transition (DT) data transfers, the REQ/ACK Offset is the maximum number of REQ assertions allowed to be outstanding before a corresponding ACK assertion is received at the target: Data transfer size is 2 bytes.
AIP	Asynchronous Information Protection. If = 1, drive supports AIP. If = 0, device does not support AIP.
AIPD	Asynchronous Information Protection Generation Disabled. If = 1, AIP generation is disabled. If = 0, AIP generation is enabled.
AIPCE	Asynchronous Information Protection Checking Enabled. If = 1, AIP protection is enabled. If = 0, AIP protection is disabled.
TWE	Transfer Width Exponent. The transfer width is 2 ⁿ bytes. If set to 0, 1 byte is selected; if set to 1, 2 bytes are selected.
PARL	Parallel Protocol. When = 1, this field indicates the negotiation occurred with a Parallel Protocol message format. When = 0, this field indicates the negotiation occurred with a Synchronous Data Request message format.
FAIR	Fairness. When set = 1, the Fairness algorithm is enabled. When set = 0, the algorithm is disabled.
SYNC	Synchronous Mode. When set = 1, the data transfer bus is operating in synchronous mode. When set = 0, the data transfer bus is operating in asynchronous mode.
QAS (Note that the value of this field must be 0 if the value in PARL is 0)	Quick Arbitrate Support Enable. When = 0, indicates that QAS is to be disabled when received from initiator and that QAS is disabled when received from target. When = 1, indicates that QAS is to be enabled when received from initiator and that QAS is enabled when received from target. See Appendix C.
DT (Note that the value of this field must be 0 if the value in PARL is 0)	Double Transition Enable. When = 0, indicates that DT DATA is to be disabled when received from initiator and that DT DATA phases are disabled when received from target. When = 1, indicates that DT DATA phases are to be enabled when received from initiator and that DT DATA phases are enabled when received from target. See Appendix C.
IU (Note that the value of this field must be 0 if the value in PARL is 0)	Information Units Enable. When = 0, indicates that information unit transfers shall not be used when received from the initiator and that information unit transfers are disabled when received from the target. When = 1, indicates that information unit transfers shall be used when received from the initiator and that information unit transfers are enabled when received from the target. See Appendix C.

*INQUIRY Command (12h) (continued)***5.11.2.9 Command Support Data Pages**

An application client can request command support data by setting the CmdDt bit of the INQUIRY command to 1, and specifying the SCSI operation code of the Command Descriptor Block (CDB) for which it wants information.

Format of the command support data and definitions of the fields follow.

Bit	7	6	5	4	3	2	1	0
Byte								
0	Peripheral Qualifier = 0			Peripheral Device Type = 0 (Direct Access Device)				
1	Reserved					Support		
2	ISO Version		ECMA Version			ANSI-Approved Version		
3 – 4	Reserved							
5	CDB Size (m – 5)							
6 – m	(MSB) CDB Usage Data (LSB)							

Figure 5–19 Command Support Data Page — Data Format

INQUIRY Command (12h) (continued)

Table 5-18 Command Support Data Page — Field Descriptions

Field	Description																		
Support	The value of the Support field describes the type of support that the disk drive provides for Command Support Data. <table border="1" data-bbox="552 399 1343 913"> <thead> <tr> <th>Value</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>Data about the requested SCSI operation code is not currently available. In this case, all data after Byte 1 is undefined.</td> </tr> <tr> <td>001b</td> <td>The device does not support the SCSI operation code requested. In this case, all data after Byte 1 is undefined.</td> </tr> <tr> <td>010b</td> <td>Reserved</td> </tr> <tr> <td>011b</td> <td>The device supports the SCSI operation code in conformance with the SCSI standard.</td> </tr> <tr> <td>100b</td> <td>Vendor-Specific</td> </tr> <tr> <td>101b</td> <td>The device supports the SCSI operation code, but in a vendor-specific manner</td> </tr> <tr> <td>110b</td> <td>Vendor-Specific</td> </tr> <tr> <td>111b</td> <td>Reserved</td> </tr> </tbody> </table>	Value	Description	000b	Data about the requested SCSI operation code is not currently available. In this case, all data after Byte 1 is undefined.	001b	The device does not support the SCSI operation code requested. In this case, all data after Byte 1 is undefined.	010b	Reserved	011b	The device supports the SCSI operation code in conformance with the SCSI standard.	100b	Vendor-Specific	101b	The device supports the SCSI operation code, but in a vendor-specific manner	110b	Vendor-Specific	111b	Reserved
Value	Description																		
000b	Data about the requested SCSI operation code is not currently available. In this case, all data after Byte 1 is undefined.																		
001b	The device does not support the SCSI operation code requested. In this case, all data after Byte 1 is undefined.																		
010b	Reserved																		
011b	The device supports the SCSI operation code in conformance with the SCSI standard.																		
100b	Vendor-Specific																		
101b	The device supports the SCSI operation code, but in a vendor-specific manner																		
110b	Vendor-Specific																		
111b	Reserved																		
ISO-Version	Must be 0.																		
ECMA-Version	Must be 0.																		
ANSI-Approved Version	Minimum operating definition for supported command.																		
CDB Size	This field contains the number of bytes in the CDB for the Operation Code being requested and the size of the CDB Usage Data in the data that is returned in response to the INQUIRY.																		
CDB Usage Data	This field contains information about the CDB for the Operation Code being queried. Note that the first byte of the CDB Usage Data contains the OpCode for the operation specified. All of the other bytes of the CDB Usage Data contain a map for bits in the CDB of the OpCode specified.																		

NOTE

The bits in the map have a 1-to-1 correspondence to the CDB for the OpCode being queried. That is, if the device senses a bit as the entire field or as part of the field of the operation, the map in CDB Usage Data contains a 1 in the corresponding bit position. If the device ignores a bit or declares a bit as "reserved" in the CDB for the OpCode being queried, the map has a 0 in that corresponding bit position. Refer to Table 5-18 for a list of the data returned for each of the OpCode values that can be sent in the INQUIRY command.

INQUIRY Command (12h) (continued)

Table 5-19 Command Support Data Page Command or Operation Codes

OpCode	Command	Hex Data Returned When INQUIRY is Received and CmdDt Bit = 1			
		Support	ANSI Version	CDB Size	CDB Usage Data
00h	TEST UNIT READY	03	01	06	00E000000003
01h	REZERO UNIT	03	01	06	01E000000003
03h	REQUEST SENSE	03	01	06	03E00000FF03
04h	FORMAT UNIT	03	01	06	04FFFFFFFF03
07h	REASSIGN BLOCKS	03	01	06	070E00000003
08h	READ (6)	03	01	06	08FFFFFFFF03
0Ah	WRITE (6)	03	01	06	0AFFFFFFFF03
0Bh	SEEK (6)	03	01	06	0BFFFFFFFF0003
12h	INQUIRY	03	01	06	12E3FF00FF03
15h	MODE SELECT (6)	03	01	06	15F10000FF03
16h	RESERVE (6)	03	01	06	16FEFFFFFFFF03
17h	RELEASE (6)	03	01	06	17FEFF000003
1Ah	MODE SENSE (6)	03	01	06	1AE8FF00FF03
1Bh	START STOP UNIT	03	01	06	1BE100000103
1Ch	REC. DIAG. RESULTS	03	01	06	1CE1FFFFFFFF03
1Dh	END DIAGNOSTIC	03	01	06	1DF700FFFF03
25h	READ CAPACITY	03	01	0A	25E0FFFFFFFFF00000103
28h	READ (10)	03	01	0A	28F8FFFFFFFFF00FFFF03
2Ah	WRITE (10)	03	01	0A	2AF8FFFFFFFFF00FFFF03
2Bh	SEEK (10)	03	01	0A	2BE0FFFFFFFFF00000003
2Eh	WRITE AND VERIFY	03	01	0A	2EF2FFFFFFFFF00FFFF03
2Fh	VERIFY	03	01	0A	2FF2FFFFFFFFF00FFFF03
35h	SYNCH. CACHE	03	01	0A	35E0FFFFFFFFF00FFFF03
37h	READ DEFECT DATA	03	01	0A	37E01F0000000000FFFF03
3Bh	WRITE BUFFER	03	01	0A	3BEFFFFFFFFE00FFFF03
3Ch	READ BUFFER	03	01	0A	3CEFFFFFFFFFFFFFFF03

(continued)

INQUIRY Command (12h) (continued)

Table 5-19 Command Support Data Page Command or Operation Codes (continued)

		Hex Data Returned When INQUIRY is Received and CmdDt Bit = 1			
Op Code	Command	Support	ANSI Version	CDB Size	CDB Usage Data
3Eh	READ LONG	03	01	0A	3EE2FFFFFFFF00FFF03
3Fh	WRITE LONG	03	01	0A	3FE0FFFFFFFF00FFF03
40h	CHANGE DEF.	03	01	0A	40E0017F000000000003
41h	WRITE SAME	03	01	0A	41E2FFFFFFFF00FFF03
4Ch	LOG SELECT	03	01	0A	4CE3C000000000FFF03
4Dh	LOG SENSE	03	01	0A	4DE1FF00000000FFF03
50h	XDWRITE (10)	03	03	0A	501CFFFFFFFF00FFF03
51h	XPWRITE (10)	03	03	0A	5118FFFFFFFF00FFF03
52h	XDREAD (10)	03	03	0A	5200FFFFFFFF00FFF03
55h	MODE SELECT (10)	03	01	0A	55100000000000FFF03
56h	RESERVE (10)	03	03	0A	5610FFFF000000FFF03
57h	RELEASE (10)	03	03	0A	5710FFFF000000FFF03
5Ah	MODE SENSE (10)	03	01	0A	5AE8FF00000000FFF03
5Eh	PERSIST. RES. IN	03	03	0A	5E1F0000000000FFF03
5Fh	PERSIST. RES. OUT	03	03	0A	5F1FFF00000000FFF03
A0h	REPORT LUNS	03	03	0C	A00000000000FFFFFFFF003
E8h	SKIP READ (10)	03	01	0A	E8E0FFFFFFFFFFFFFF03
EAh	SKIP WRITE (10)	03	01	0A	EAE0FFFFFFFFFFFFFF03

5.12 LOG SELECT Command (4Ch)

The drive collects and stores performance data and error summaries in counters. The LOG SELECT command is used to zero these counters. The LOG SELECT command is a complementary command to the LOG SENSE command

Bit	7	6	5	4	3	2	1	0
Byte								
0	Operation Code (4Ch)							
1	Reserved						PCR	SP
2	PC		Reserved					
3 – 6	Reserved							
7 – 8	Parameter List Length							
9	Control							

Figure 5-20 LOG SELECT Command Descriptor Block — Data Format

Table 5-20 LOG SELECT Command Descriptor Block — Field Descriptions

Field	Description
PCR	Parameter Code Reset. Must be 1. Causes all the implemented counters to be set to 0.
SP	Save Parameters. A value of 1 indicates that certain counters are to be cleared from non-volatile memory.
PC	Page Control. Must be 01b OR 11b.
Parameter List Length	Must be 0.

5.13 LOG SENSE Command (4Dh)

NOTE

Log Sense data pages require special interpretation and also are subject to change. For assistance with the Log Sense data pages, contact your Quantum Applications Engineer.

The drive collects operational information and stores these statistics as *log data*. Log data are grouped by category into *log pages*. The LOG SENSE command allows an initiator to retrieve the stored log data. The LOG SENSE command is a complementary command to the LOG SELECT command.

Each log page contains one or more pieces of information. Each piece of information is referred to as a *parameter*. There are two types of parameters: *values* and *lists*. In general, error and performance counters are reported as values. For example, the count of “Hardware Idled Due To Invalid Command” is returned as a 32-bit integer. Events are reported as lists. For example, the last known failure might be reported with the ASCII string “Last known failure code 960A(X)”.

Table 5–21 lists the log pages supported by the drive. Contact your Quantum Applications Engineer for more information.

Table 5-21 Disk Drive Log Pages

Page Code	Description
00h	Supported Log Pages
01h	Buffer Overruns and Underruns
02h	Write Error Counter
03h	Read Error Counter
05h	Verify Error Counter
06h	Non-Medium Error Counter Page
07h	Last n-Error Events Page
08h	Format Status Page
0Fh	Application Client Page
10h	Self Test Results Page
2Fh ¹	SMART Page
31h	Last 500 Errors
36h ¹	Early Warning System (EWS) Page
37h ¹	Seek Performance Summary
38h ¹	Servo Events Counter
39h ¹	Bad Block Replacement Summary
3Ah ¹	Disk Error Recovery Counters Page
3Bh ¹	DER Description
3Dh ¹	ECC Counters and Summary Page
3Eh ¹	SCSI Bus Events Page
3Fh	ECC On The Fly

NOTE: Page codes 2Fh, 31h, 36h, 37h, 38h, 39h, 3Ah, 3Bh, 3Dh, 3Eh, and 3Fh are vendor-specific pages that are unique to Quantum drives. Contact your Quantum representative for more information.

LOG SENSE Command (4Dh) (continued)

5.13.1 LOG SENSE Command Descriptor Block

The Command Descriptor Block for the LOG SENSE command is shown in Figure 5-21. Table 5-22 contains field descriptions.

Bit	7	6	5	4	3	2	1	0
Byte								
0	Operation Code (4Dh)							
1	Reserved						PPC	SP
2	PC		Page Code					
3 - 4	Reserved							
5 - 6	Parameter Pointer							
7 - 8	Allocation Length							
9	Control							

Figure 5-21 LOG SENSE Command Descriptor Block — Data Format

Table 5-22 LOG SENSE Command Descriptor Block — Field Descriptions

Field	Description
PPC	The Parameter Pointer Control bit controls the type of parameters that can be requested from the drive. This bit must be 0, indicating that all log parameters for the specified page (subject to the allocation length specified) are returned to the initiator.
SP	The Save Pages bit specifies whether the parameters are to be saved. This bit must be 0, indicating that no parameters are to be saved and are reset at power-on or by a TARGET RESET.
PC	The Page Control field defines the type of parameter values to be selected. The field must be 01b indicating that the current values are to be returned. Mode 11b (return default values) is not supported because all counters have a default value of 0.
Page Code	The value specified as the Page Code determines the page to be returned. Table 5-21 contains a list of supported log pages and their page codes.
Parameter Pointer	This field is related to the PPC field. This bit must be 0.
Allocation Length	This specifies the number of bytes of data that the drive is allowed to pass during the DATA IN phase. The requested page is truncated if its length exceeds the number of bytes specified in this field.

LOG SENSE Command (4Dh) (continued)

5.13.2 LOG SENSE Log Pages

The log pages that are returned from the drive have a common format that is shown in Figure 5–22. Each page contains a 4-byte header followed by one or more log parameters. Refer to Table 5–23.

Bit Byte	7	6	5	4	3	2	1	0
0	Reserved		Page Code					
1	Reserved							
2 – 3	Page Length (n-3)							
4 to x+3	Log Parameter (first) (length x) ¹							
•	•							
•	•							
•	•							
•	•							
n-y+1 to n	Log Parameter (last) (length y) ¹							

NOTE: ¹ Length x or y is the sum of parameter header and parameter data value bytes.

Figure 5-22 LOG SENSE Log Page Format - Data Format

Table 5-23 LOG SENSE Log Page Format - Field Descriptions

Field	Description
Page Code	The page code as given in Table 5–20.
Page Length	The allocation length for the page minus the 4-byte header.
Log Parameter	One or more log data entities that are returned as par of a Log Page. Each Log Parameter starts with a 4-byte header followed by one or more bytes of <i>value</i> data. In most cases, the parameter value is a 4-byte longword that contains the present value of an error or performance counter. See Figure 5–23 for the format of a generic log parameter.

LOG SENSE Command (4Dh) (continued)

Bit	7	6	5	4	3	2	1	0
Byte								
0 - 1	Parameter Code							
2	DU	DS	TSD	ETC		TMC		LP
3	Page Length (n-3)							
4 - n	Parameter Value							

Figure 5-23 Generic Log Parameter - Data Format

Table 5-24 Generic Log Parameter - Field Descriptions

Field	Description
Parameter Code	A code which uniquely identifies each parameter on a given Log Page. For example, the code 8002h on the Seek Performance Summary page reports the average seek time while the same code on the Block Replacement Summary page reports the number of blocks replaced. The following pages in this manual contain the lists of applicable parameters for each Log page.
DU	Disable Update. This parameter is 1 when updates are not enabled.
DS	Disable Save. This parameter is 1 when parameters are not saved by the drive.
TSD	Target Save Disable. This parameter is 0 indicating that the drive provides a target-specific way of saving parameters.
ETC	Enable Threshold Comparison. This parameter is always 0. The drive does not use thresholds for any of its parameters.
TMC	Threshold Met Comparison. This parameter is always 0. The drive does not use thresholds for any of its parameters.
LP	List Parameter. When 0, indicates that the parameter value for this parameter is a numeric <i>value</i> . When LP is a 1, it indicates that the parameter value is an alphanumeric ASCII-string <i>list</i> . This parameter is always 0.
Parameter Length	Specifies the length, in bytes, of the parameter's value.
Parameter Value	Contains the parameter's current value when the PC field of the LOG SENSE command is 01b.

NOTE: The DU, DS, TSD, ETC, TMC and LP fields are collectively referred to as the Parameter Control Byte. This byte generally has a value of 0; however, if the page is non-volatile but the drive is not spun up or the GLTSD bit is set in Mode Page 0Ah, then the value would be 20h. If the page is volatile, then the value is 60h.

5.14 MODE SELECT (6) Command (15h)

SCSI refers to the drive's operational parameters as its *mode parameters*. SCSI groups the mode parameters by function into a set of data structures referred to as *mode pages*. The MODE SELECT (6) command allows the initiator to modify some of these mode pages and thereby control some of the drive's operational characteristics. The Save Page (SP) option in the Command Descriptor Block makes the changes permanent. The new mode parameters are then stored in the drive's non-volatile memory.

Bit	7	6	5	4	3	2	1	0
Byte								
0	Operation Code (15h)							
1	Reserved			PF	Reserved			SP
2 – 3	Reserved							
4	Parameter Length List							
5	Control							

Figure 5-24 MODE SELECT (6) Command Descriptor Block — Data Format

Table 5-25 MODE SELECT (6) Command — Field Descriptions

Data Field	Description
PF	Page Format. The drive ignores the content of this field and produces a mode parameter list that contains a mode header optionally followed by a sector descriptor and one or more pages that conform to the format shown here.
SP	Save Pages. When SP = 0, the drive performs the specified MODE SELECT operation and does not save any pages. When SP = 1, the drive performs the specified operation and stores all savable pages, (PS = 1 on MODE SENSE return) including any sent during the data-out buffer transfer. The changes made when SP = 1 become permanent changes to the drive's SCSI setup.
Parameter List Length	This field tells the drive how many bytes of Mode Parameters to fetch in the data-out buffer transfer. A CHECK CONDITION status is returned if this value truncates a page.

*MODE SELECT (6) Command (15h) (continued)***5.14.1 Initiator-Changeable Mode Pages**

Table 5-26 lists the mode pages that are supported by the drive. An initiator can change these pages by supplying them, with the desired changes included, in the MODE SELECT command's data-out buffer transfer.

The initiator should first use a MODE SENSE command to read the appropriate pages and leave non-changeable values as read when the initiator subsequently writes the changeable pages with the MODE SELECT command. If a non-changeable field contains an invalid value, the drive returns a CHECK CONDITION status.

Table 5-26 Initiator-Changeable Mode Pages

Page Code	Page Name	Function	Size In Bytes
00h	Unit Attention Control Page	Unit Attention reporting (enables or disables)	4
01h	Read-Write Error Recovery Page	Medium Access Error recovery and reporting procedures for READ and WRITE commands	12
02h	Disconnect-Reconnect Page	Bus behavior during data transfers	16
07h ¹	Verify Error Recovery Page	Medium Access Error recovery and reporting procedures for the VERIFY command	12
08h ¹	Caching Page	Cache policy	20
0Ah ¹	Control Mode Page	Command processing policy	12
0Ch	Notch and Partition Page ²	Drive geometry reporting	24
10h	XOR Control Page	Controls XOR operating parameters	24
1Ah	Power Condition Page	Enable and set time intervals for Idle and Standby modes	12
1Ch ¹	Information Exceptions Control Page	Defines methods to control reporting and operations of specific informational exception conditions.	12
39h	Quantum-Unique Page		22

NOTES:

1. This page is not supported by the SCSI-1/CCS Operating Mode.
2. The Notch and Partition page is changeable but NOT savable.

MODE SELECT (6) Command (15h) (continued)

5.14.2 Mode Page Types

The drive maintains three distinct sets of mode pages. They are the *current page*, the *default page*, and the *saved page*. The drive also reports a fourth set of *changeable* pages. The page types are defined in Table 5-27.

Table 5-27 Mode Page Types

Page Type	Definition
Current	<p>The <i>current</i> mode page set applies to all <i>initiators</i> and defines the drive's mode. The SCSI-2 specification states that a drive can maintain Mode parameters on a per-initiator basis if it so chooses. The Drives do not support this option.</p> <p>The <i>current</i> mode page set contains the values supplied in the last MODE SELECT command received from an initiator. If no initiator has sent a MODE SELECT command since the drive was last reset or powered up, the current pages contain</p> <ul style="list-style-type: none"> • <i>Saved</i> values if saved pages exist (from a previous MODE SELECT command SP parameter). • <i>Default</i> values if pages have never been saved. <p>The drives generate a Unit Attention condition for all initiators (except for the one that was the source of the MODE SELECT command) whenever one initiator modifies the Mode parameters.</p>
Default	The <i>default</i> mode page set contains the factory default values that are listed in each page's description.
Saved	The <i>saved</i> mode page set contains values preserved in the drive's non-volatile memory by a previous SP-modified MODE SELECT command.
Changeable	The <i>changeable</i> mode page set provides a means for an initiator to determine which pages it is allowed to change and the specific bits within those pages that it is allowed to change. This page set is read-only and is fetched with a MODE SENSE command.

5.14.3 Mode Parameter List

Figure 5-25 shows the format of the Mode Parameter List that is passed by the initiator to the drive during the command's data-out buffer transfer. Table 5-28 describes the fields. Figure 5-26 and Table 5-29 provides a description of the data format and fields of the Mode Parameter Header.

Bit	7	6	5	4	3	2	1	0
Byte								
0 – 3	Mode Parameter Header							
4 – 11	Block Descriptor (optional)							
4 – n or 12 – n	Page(s) (optional)							

Figure 5-25 Mode Parameter List — Data Format

MODE SELECT (6) Command (15h) (continued)

Table 5-28 Mode Parameter List — Field Descriptions

Field	Description
Mode Parameter Header	Contains information about the remainder of the parameter list and is always present (see Figure 5-26 and Table 5-29).
Block Descriptor	Allows the initiator to set the drive's Logical Block Size and number of Logical Block Addresses (see Figure 5-27 and Table 5-30).
Page(s)	The page code(s) of the pages that are a part of this command.

Bit Byte	7	6	5	4	3	2	1	0
0	Mode Data Length							
1	Medium Type							
2	Device-Specific Parameter							
3	Block Descriptor Length							

Figure 5-26 Mode Parameter Header (6-Byte) — Data Format

Table 5-29 Mode Parameter Header — Field Descriptions

Field	Description
Mode Data Length	Reserved. Must be 0.
Medium Type	Ignored by the drive.
Device-specific Parameter	Ignored by the drive.
Block Descriptor Length	Zero (0) if no Block Descriptor is supplied. The length is eight (8) if a Block Descriptor is supplied.

Bit Byte	7	6	5	4	3	2	1	0
0	Density Code							
1 – 3	Number of Blocks							
4	Reserved							
5 – 7	Block Length							

Figure 5-27 Mode Parameter Block Descriptor – Data Format

*MODE SELECT (6) Command (15h) (continued)**Table 5-30 Mode Parameter Block Descriptor — Field Descriptions*

Field	Description
Density Code	Ignored by the drive.
Number of Blocks	This field specifies the number of logical blocks on the unit that are affected by the current MODE SELECT command. When this field is 0, the MODE SELECT parameters are put into effect for all logical blocks on the unit. The maximum value allowed is the total number of initiator-addressable logical blocks for the value specified in the Block Length field. A value that is less than the maximum value will decrease the number of initiator-accessible logical blocks to the amount specified.
Block Length	This field specifies the length, in bytes, of each logical sector.

5.14.3.1 Categories of Changeable Pages

The drive's changeable pages are described on the following pages. The data fields for each of these pages fall into one of the categories described in Table 5-31.

Table 5-31 Categories of Changeable Pages

Category	Description
Fully Supported	A field in this category can be modified by the MODE SELECT command, can be saved, and can affect the drive's processing. The drive uses the value in the field and, if the SP bit is set, preserves the new value of the field in non-volatile memory on the drive. (The PS bit must have been set when the mode page was returned with a MODE SENSE command.) Fully Supported fields are <u>underlined</u> in the Mode Parameter Pages shown on the following pages. The PS bit is underlined (meaning it will be set on a MODE SENSE command) only if the page is savable.
Temporary	A field in this category is fully supported, with one exception; the drive does not save the field contents. The drive does not give an error indication if an attempt is made to save a temporary field. Temporary fields are also underlined in the Mode Parameter Page illustrations but the PS bit is not.
Ignored	A value in this field is never used or validated; it is never looked at by the drive. Ignored fields are not underlined in the Mode Parameter Pages' figures or tables and are not described in the Field Description tables.

MODE SELECT (6) COMMAND (15h) (continued)

5.14.3.2 Unit Attention Condition Page (00h)

The Unit Attention Page is used to enable or disable the generation of Unit Attention Conditions.

Bit	7	6	5	4	3	2	1	0
Byte								
0	<u>PS</u>	Rsvd	Page Code (00h)					
1	Page Length (02h)							
2	Reserved			<u>DUA</u>	Reserved			
3	Reserved (00h)							

Figure 5-28 Unit Attention Condition Page — Data Format

Table 5-32 Unit Attention Condition Page — Data Format

Field	Default Value	Description
<u>PS</u>	1	Parameters Savable. This bit is only used with the MODE SENSE command. A returned value of 1 indicates that the target is capable of saving the page in a non-volatile, vendor-specific location. The bit is reserved with the MODE SELECT command.
<u>DUA</u>	0	Disable Unit Attention. When set to the default value of 0, the drive reports a CHECK CONDITION for any UNIT ATTENTION condition (power-on, SCSI bus reset, etc.) If set to 1, the drive does not report any UNIT ATTENTION conditions.

MODE SELECT (6) Command (15h) (continued)

5.14.3.3 Read-Write Error Recovery Page (01h)

The Read-Write Error Recovery page controls the drive's response to error conditions that arise during the course of READ (6), READ (10), WRITE (6), WRITE (10) command processing and during the write portion of WRITE AND VERIFY command processing.

Bit Byte	7	6	5	4	3	2	1	0
0	<u>PS</u>	Rsvd	Page Code (01h)					
1	Page Length (0Ah)							
2	<u>AWRE</u>	<u>ARRE</u>	<u>TB</u>	<u>RC</u>	EER	<u>PER</u>	DTE	<u>DCR</u>
3	<u>Read Retry Count</u>							
4	<u>Correction Span</u>							
5	Head Offset Count							
6	Data Strobe Offset Count							
7	Reserved							
8	<u>Write Retry Count</u>							
9	Reserved							
10 – 11	<u>Recovery Time Limit</u>							

NOTE: Fully supported fields are underlined. Unsupported and ignored fields are not underlined. The PS bit is underlined, meaning it will be set on a MODE SENSE command, since the page is savable.

Figure 5-29 Read-Write Error Recovery Page — Data Format

MODE SELECT (6) Command (15h) (continued)

Table 5-33 Read-Write Error Recovery Page — Field Descriptions

Field	Default Value	Description
<u>PS</u>		1 Parameters Savable. This bit is only used with the MODE SENSE command. A returned value of 1 indicates that the target is capable of saving the page in a non-volatile, vendor-specific location. The bit is reserved with the MODE SELECT command.
<u>AWRE</u>	1	Automatic Write Reallocation Enable. When the value is 1, it causes the drive to automatically reallocate bad sectors encountered during a write operation. If it succeeds in replacing the bad sector, the drive adds the revector sector's Logical Block Number to the Grown Defect List. Reallocation is done only when the drive is able to recover the sector's data unless the Reallocate Uncorrected Errors bit of the Quantum (Vendor) Special Function Control Page is set. When AWRE is 0, the drive does not automatically reallocate bad sectors encountered during a write operation. See also Quantum-Unique Page 39h.
<u>ARRE</u>	1	Automatic Read Reallocation Enable. When the value is 1, it causes the drive to automatically reallocate bad sectors encountered during a read operation. If it succeeds in replacing the bad sector, the drive adds the revector sector's Logical Block Number to the Grown Defect List. Reallocation is done only when the drive is able to recover the sector's data unless the HrdBBR bit of the Quantum (Vendor) Special Function Control Page is set. When ARRE is 0, the drive does not automatically reallocate bad sectors encountered during a read operation. See also Quantum-Unique Page 39h (RUEE).
<u>TB</u>		0 Transfer Block. When the value is 1, the drive sends best-guess information to the initiator when it encounters a defective sector whose data cannot be fully recovered. When the value is 0, data that cannot be fully recovered is not returned to the initiator.
<u>RC</u>	0	Read Continuous. When the value is 0, the drive attempts error recovery when it encounters an error. When the value is 1, the drive suppresses error recovery. The drive ignores any value placed in this field.
<u>EER</u>	0	Enable Early Recovery. Not supported by the drive.
<u>PER</u>	0	Post Error Recovery. When the value is 0, the drive returns a GOOD status in cases where error recovery is invoked and successful. When the value is 1, the drive returns a CHECK CONDITION status for recovered errors, sets the sense key to RECOVERED ERROR, and posts the appropriate additional sense key.
<u>DTE</u>	0	Disable Transfer on Error. Not supported by the drive.
<u>DCR</u>	0	Disable Correction. When the value is 0, ECC correction is applied as a means of recovering erred data. On write operations, a valid ECC is written to each data sector regardless of the value in the DCR bit. When the value is 1, ECC correction is disabled (but not retries or reads), and ECC errors cannot be corrected. However, on-the-fly is always enabled.

MODE SELECT (6) Command (15h) (continued)

Table 5-33 Read-Write Error Recovery Page — Field Descriptions (continued)

Field	Default Value	Description
<u>Read Retry Count</u>	4	Specifies the number of retries attempted for a failed read operation.
<u>Correction Span</u>	170	Specifies the maximum number of bits in a data sector that the drive is allowed to correct. A value less than the default value can cause an otherwise correctable error to be reported as uncorrectable. If the initiator specifies a value larger than the default value, the drive automatically rounds down the number to the default value. A value of 0 causes the drive to use the default value. The drive ignores any value placed in this field.
Head Offset Count	0	Not supported by the drive.
Data Strobe Offset Count	0	Not supported by the drive.
<u>Write Retry Count</u>	8	Specifies the number of retries attempted for a failed write operation.
<u>Recovery Time Limit</u>	0	Specifies the maximum time, in milliseconds, that a retry is attempted on a failed sector during a read or write operation. When the value is 0, it means that there is no time limit. The minimum permissible value, however, is 100 milliseconds.

NOTE: Fully supported and temporary fields are underlined. Unsupported and ignored fields are not underlined. The PS bit is underlined, meaning it will be set on a MODE SENSE command, since the page is savable

MODE SELECT (6) Command (15h) (continued)

5.14.3.4 Disconnect-Reconnect Page (02h)

The Disconnect-Reconnect Page provides the application client the means to attempt to optimize the performance of the delivery subsystem.

Bit Byte	7	6	5	4	3	2	1	0
0	<u>PS</u>	Rsvd	Page Code (02h)					
1	Page Length (0Eh)							
2	<u>Buffer Full Ratio</u>							
3	<u>Buffer Empty Ratio</u>							
4 – 5	Bus Inactivity Limit							
6 – 7	<u>Disconnect Time Limit</u>							
8 – 9	Connect Time Limit							
10 – 11	<u>Maximum Burst Size</u>							
12	EMDP	<u>Fairness Arbitration</u>			<u>DImm</u>	<u>DTDC</u>		
13	Reserved							
14 – 15	First Burst Size							

NOTE: Fully supported fields are underlined. Unsupported and ignored fields are not underlined. The PS bit is underlined, meaning it will be set on a MODE SENSE command, since the page is savable.

Figure 5-30 Disconnect-Reconnect Page — Data Format

MODE SELECT (6) Command (15h) (continued)

Table 5-34 Disconnect-Reconnect) — Field Descriptions

Field	Default Value	Description						
<u>PS</u>	0	Parameters Savable. This bit is only used with the MODE SENSE command. A returned value of 1 indicates that the target is capable of saving the page in a non-volatile, vendor-specific location. The bit is reserved with the MODE SELECT command.						
<u>Buffer Full Ratio</u>	0	The Buffer Full Ratio indicates how full the buffer should be prior to requesting an interconnect tenancy. Applies to READ operations only. The default of 0 indicates to the drive that it should use its default ratio.						
<u>Buffer Empty Ratio</u>	0	The Buffer Empty Ratio indicates how full the buffer should be prior to requesting an interconnect tenancy. Applies to READ operations only. The default of 0 indicates to the drive that it should use its default ratio.						
Bus Inactivity Limit	0	The Bus Inactivity Limit indicates the maximum duration of any interconnect tenancy during which no data is transferred. Not supported by the drive.						
<u>Disconnect Time Limit</u>	0	The Disconnect Time Limit specifies the minimum time, in 100 microsecond increments, that the drive waits after releasing the SCSI bus before reselecting. The drive supports a maximum value of 0xFF. Values greater than 0xFF are treated as 0xFF. The default value of 0 indicates that the drive can reslect immediately after releasing the bus.						
Connect Time Limit	0	The Connect Time Limit indicates the maximum duration of a single interconnect tenancy. Not supported by the drive.						
<u>Maximum Burst Size</u>	0	The Maximum Burst Size specifies the maximum amount of data that the drive will transfer during a data phase before disconnecting from the bus, assuming it was given the disconnect privilege. This value is expressed in increments of data sectors (that is, a value of 1 indicates 512 bytes, a value of 2 indicates 1024 bytes). The default value of 0 indicates that there is no limit on the amount of data transferred per connection. This field is reserved for the SCSI-1/CCS operating mode.						
EMDP	0	Enable MODIFY DATA POINTER. Not supported by the drive.						
<u>Fairness Arbitration</u>	111b	Indicates whether or not the target should use fair or unfair (e.g., priority) arbitration when beginning the interconnect tenancy. <table border="1"> <thead> <tr> <th>Value</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>Disables the fairness algorithm. A fixed priority scheme based on the SCSI ID assigned is used.</td> </tr> <tr> <td>Non-Zero</td> <td>Any non-zero value enables the fairness algorithm (SCSI SPI-2 ANSI standard, Appendix C).</td> </tr> </tbody> </table>	Value	Description	000b	Disables the fairness algorithm. A fixed priority scheme based on the SCSI ID assigned is used.	Non-Zero	Any non-zero value enables the fairness algorithm (SCSI SPI-2 ANSI standard, Appendix C).
Value	Description							
000b	Disables the fairness algorithm. A fixed priority scheme based on the SCSI ID assigned is used.							
Non-Zero	Any non-zero value enables the fairness algorithm (SCSI SPI-2 ANSI standard, Appendix C).							
<u>Dlmm</u>	0	Disconnect Immediate. When set to a 1, to reduce on-bus time, the drive is forced to disconnect from the SCSI Bus after a SCSI Command is received. Disconnect Immediate works only when a typical, error-free command is received. Disconnect Immediate does not apply in exception conditions (for example, SDTR received, error, etc.).						

MODE SELECT (6) Command (15h) (continued)

Table 5-34 Disconnect-Reconnect Page — Field Descriptions (continued)

<u>Field</u>	<u>Default Value</u>	<u>Description</u>										
<u>DTDC</u>		Data Transfer Disconnect Control. This field of three bits defines further restrictions when a disconnect is permitted. These are as follows: <table border="1"> <thead> <tr> <th><u>Value</u></th> <th><u>Description</u></th> </tr> </thead> <tbody> <tr> <td>000b</td> <td>Data transfer disconnect control is not used. Disconnect is controlled by other fields in this page.</td> </tr> <tr> <td>001b</td> <td>A target shall not attempt to disconnect once the data transfer of a command has started until all data of the command has been transferred. The connect time limit and bus inactivity limit are ignored during the transfer.</td> </tr> <tr> <td>010b</td> <td>Reserved</td> </tr> <tr> <td>011b</td> <td>A target shall not attempt to disconnect once the data transfer of a command has started, until the command is complete. The connect time limit and bus inactivity limit are ignored once the transfer has started.</td> </tr> </tbody> </table>	<u>Value</u>	<u>Description</u>	000b	Data transfer disconnect control is not used. Disconnect is controlled by other fields in this page.	001b	A target shall not attempt to disconnect once the data transfer of a command has started until all data of the command has been transferred. The connect time limit and bus inactivity limit are ignored during the transfer.	010b	Reserved	011b	A target shall not attempt to disconnect once the data transfer of a command has started, until the command is complete. The connect time limit and bus inactivity limit are ignored once the transfer has started.
<u>Value</u>	<u>Description</u>											
000b	Data transfer disconnect control is not used. Disconnect is controlled by other fields in this page.											
001b	A target shall not attempt to disconnect once the data transfer of a command has started until all data of the command has been transferred. The connect time limit and bus inactivity limit are ignored during the transfer.											
010b	Reserved											
011b	A target shall not attempt to disconnect once the data transfer of a command has started, until the command is complete. The connect time limit and bus inactivity limit are ignored once the transfer has started.											
First Burst Size	0	This field indicates the maximum amount of data that a target may transfer for a command during the same interconnect in which it receives the command.										

NOTE: Fully supported and temporary fields are underlined. Unsupported fields are not underlined. The PS bit is underlined, meaning it will be set on a MODE SENSE command, since the page is savable.

*MODE SELECT (6) Command (15h) (continued)***5.14.3.5 Verify Error Recovery Page (07h)**

The Verify Error Recovery Page controls the drive's response to error conditions that arise during the VERIFY command and during the verify operation of the WRITE AND VERIFY command. The Verify Error Recovery Page is not supported in the SCSI-1/CCS Operating Mode.

Bit	7	6	5	4	3	2	1	0
Byte								
0	<u>PS</u>	Rsvd	Page Code (07h)					
1	Page Length (0Ah)							
2	Reserved				EER	<u>PER</u>	DTE	<u>DCR</u>
3	<u>Verify Retry Count</u>							
4	Verify Correction Span							
5 – 9	Reserved							
10 – 11	<u>Verify Recovery Time Limit</u>							

NOTE: Fully supported and temporary fields are underlined. Informational fields are not underlined. The PS bit is underlined, meaning it will be set on a MODE SENSE command, since the page is savable

Figure 5–31 Verify Error Recovery Page — Data Format

MODE SELECT (6) Command (15h) (continued)

Table 5-35 Verify Error Recovery Page — Field Descriptions

Field	Default Value	Description
<u>PS</u>	0	Parameters Savable. This bit is only used with the MODE SENSE command. A returned value of 1 indicates that the target is capable of saving the page in a non-volatile, vendor-specific location. The bit is reserved with the MODE SELECT command.
EER	0	Enable Early Recovery. Not supported by the drive.
<u>PER</u>	0	Post Error Recovery. When the value is 1, the drive returns CHECK CONDITION status for recovered errors, sets the sense key to RECOVERED ERROR, and posts the appropriate additional sense code. When set to the default value of 0, it causes the drive to return GOOD status in cases where error recovery is invoked and is successful. The only recoverable error that can occur during a VERIFY operation is a failed read that succeeds after retries.
DTE	0	Disable Transfer on Error. Not supported by the drive.
<u>DCR</u>	0	Disable Correction. When the value is 1, it disables ECC correction (but not retries) on the READ portion of a VERIFY operation. When set to the default value of 0, the ECC correction is always applied as a means of recovering erred data.
<u>Verify Retry Count</u>	4	The Verify Retry Count specifies the number of retries that are attempted per sector on a failed READ portion of a VERIFY operation.
Verify Correction Span	170	The Verify Correction Span specifies the maximum number of bits in a sector that the drive is allowed to correct. A value less than the default value can cause an otherwise correctable error to be reported as uncorrectable. If the initiator specifies a value larger than 80, the drive rounds down to 80 without reporting the change. A value of 0 causes the drive to use its default value.
<u>Verify Recovery Time Limit</u>	0	The Verify Recovery Time Limit specifies (in milliseconds) the maximum time that a retry is attempted on a failed sector during the read portion of a verify operation. The default value of 0 indicates that there is no time limit. The minimum permissible is 100 milliseconds.

NOTE: Fully supported and temporary fields are underlined. Unsupported fields are not underlined. The PS bit is underlined, meaning it will be set on a MODE SENSE command, since the page is savable.

MODE SELECT (6) Command (15h) (continued)

5.14.3.6 Caching Page (08h)

The Caching Page controls the drive's cache management policy.

Bit	7	6	5	4	3	2	1	0
Byte								
0	<u>PS</u>	Rsvd	Page Code (08h)					
1	Page Length (12h)							
2	Reserved					<u>WCE</u>	<u>MF</u>	<u>RCD</u>
3	Demand Read Retention Priority				Demand Write Retention Priority			
4 – 5	Disable Pre-Fetch Transfer Length							
6 – 7	<u>Minimum Pre-Fetch</u>							
8 – 9	Maximum Pre-Fetch							
10 – 11	Maximum Pre-Fetch Ceiling							
12	FSW	LBCSS	DRA	VS	VS	Reserved		
13	Number of Cache Segments							
14 – 15	Cache Segment Size							
16	Reserved							
17 – 19	Non-Cache Segment Size							

- NOTES:**
1. Fully supported and temporary fields are underlined. Informational fields are not underlined. The PS bit is underlined, meaning it will be set on a MODE SENSE command, since the page is savable
 2. Bytes 12 to 19 are returned only if the device operating definition is SCSI-3.

Figure 5-32 Caching Page — Data Format

MODE SELECT (6) Command (15h) (continued)

Table 5-36 Caching Page — Field Descriptions

Field	Default Value	Description
<u>PS</u>	1	Parameters Savable. This bit is only used with the MODE SENSE command. A returned value of 1 indicates that the target is capable of saving the page in a non-volatile, vendor-specific location. Bit reserved with MODE SELECT command.
<u>WCE</u>	1	Write Cache Enable. The default value of 1 enables write-back caching when processing a WRITE command. The drive returns a GOOD status and COMMAND COMPLETE only after successfully writing all the data to the media. When WCE = 1, write-back caching is enabled while processing a WRITE command. The drive returns a GOOD status and COMMAND COMPLETE only after fetching the data from the initiator and placing it in the drive's cache memory. Note: For maximum data integrity, it is <u>strongly recommended</u> that when WCE is set to 1, the Read-Write Error Recovery Page (01h) fields be set as follows: <ul style="list-style-type: none"> • AWRE to 1 • RC to 0 • DCR to 0 • Correction Span set to 170 • Head Offset to 0 • Servo Offset to 0
<u>MF</u>	0	Multiplication Factor (value must be 0). The drive interprets the minimum and maximum pre-fetch fields values in terms of the number of logical blocks for each of the respective types of prefetch.
<u>RCD</u>	0	Read Cache Disable. The default value of 0 enables the read cache capability and allows the drive to use cache-resident data or medium data to satisfy a READ command. When the value is 1, read caching is disabled and the data returned in response to a READ command is fetched from the medium and not from the drive's cache memory.
Demand Read Retention Priority	0	This field advises the drive on the retention priority to assign data read into the cache that has also been transferred from the drive to the initiator. Not supported by the drive.
Demand Write Retention Priority	0	This field advises the drive on the retention priority to assign data written into the cache that has also been transferred from the cache memory to the medium. Not supported by the drive.
Disable Pre-Fetch Transfer Length	FFFFh	This field specifies the selective disabling of anticipatory pre-fetch on long transfer lengths. The value in this field is compared to the number of sectors requested by the current READ command. If the number of sectors is greater than the value specified, anticipatory pre-fetch is not performed for the command. When 0, anticipatory pre-fetch is disabled. Ignored by the drive.

(continued)

MODE SELECT (6) Command (15h) (continued)

Table 5-36 Caching Page — Field Descriptions (continued)

Field	Default Value	Description
<u>Minimum Pre-fetch</u>	0	This field indicates the number of sectors that are to be pre-fetched. See also MF.
Maximum Pre-fetch	0421h	This field indicates the maximum amount of data to pre-fetch into the cache as a result of one READ command. Not supported by the drive.
Maximum Pre-fetch Ceiling	0421h	Specifies an upper limit on the number of logical blocks computed as the maximum pre-fetch. If this number of sectors is greater than the maximum pre-fetch, then the number of logical blocks to pre-fetch is reduced to the value stored in the maximum pre-fetch ceiling field. Not supported by the drive.
FSW	0	The Force Sequential Write bit. Not supported by the drive.
LBCSS	0	The Logical Block Cache Segment Size bit. Not supported by the drive.
DRA	0	The Disable Read-Ahead bit. Not supported by the drive.
VS		The Vendor-Specific bits. Not supported by the drive.
Number of Cache Segments	27	Ignored by the drive.
Cache Segment Size	0	Not supported by the drive.
Non-Cache Segment Size	0	Not supported by the drive.

NOTE: Fully supported fields are underlined. Unsupported and ignored fields are not underlined. The PS bit is underlined, meaning it will be set on a MODE SENSE command, since the page is savable.

MODE SELECT (6) Command (15h) (continued)

5.14.3.7 Control Mode Page (0Ah)

The Control Mode page controls the drive's command processing and error handling policies.

Bit	7	6	5	4	3	2	1	0
Byte								
0	<u>PS</u>	Rsvd	Page Code (0Ah)					
1	Page Length (0Ah)							
2	Reserved					<u>GLTSD</u>	RLEC	
3	<u>Queue Algorithm Modifier</u>			Reserved		<u>Qerr</u>	<u>DQue</u>	
4	<u>EECA</u>	RAC	Reserved		RAERP	UAAERP	EAERP	
5	Reserved							
6 – 7	Ready AEN Holdoff Period							
8 – 9	Busy Timeout Period							
10 – 11	Reserved							

- NOTES:**
1. Fully supported and temporary fields are underlined. Unsupported and ignored fields are not underlined. The PS bit is underlined, meaning it will be set on a MODE SENSE command, since the page is savable.
 2. Bytes 8 to 11 are returned only if the device operating definition is SCSI-3.

Figure 5-33 Control Mode Page — Data Format

MODE SELECT (6) Command (15h) (continued)

Table 5-37 Control Mode Page — Field Descriptions

Field Value	Default	Description
<u>PS</u>	1	Parameters Savable. This bit is only used with the MODE SENSE command. A returned value of 1 indicates that the target is capable of saving the page in a non-volatile, vendor-specific location. The bit is reserved with the MODE SELECT command.
<u>GLTSD</u>	0	Global Logging Target Save Disable. If this bit = 1, the ability to save log ages to disk is disabled. If a LOG SELECT with the SP bit is received by the drive, the current cumulative values will be cleared. The on-disk values will not be changed. If a LOG SENSE with the SP bit set is received by the drive and the GLTSD bit is set and the unit is ready, the last LOG counter values written to disk will be returned (not the current cumulative values). If a STOP UNIT command or a LOG SENSE command with the SP bit set is issued to the drive, the cumulative values of the Log Counters page will be written to disk. A LOG SENSE command will return the cumulative values since the counters were cleared, not since the drive was powered on.
<u>RLEC</u>	0	Report Log Exception Condition. The drive does not report log exception conditions. Not supported by the drive.
<u>Queue Algorithm Modifier</u>	0001b	This field specifies restrictions on the algorithm used for re-ordering commands that are tagged with the SIMPLE message. A value of 0 specifies that the drive orders the actual execution sequence with a SIMPLE statement. A value of 1 (Unrestricted Re-ordering Allowed) specifies that the drive can re-order the actual execution sequence of the commands with a SIMPLE in any manner. 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.
<u>QErr</u>	0	Queue Error Management. This field controls the drive's processing of queued commands when a Contingent Allegiance (CA) condition or Extended Contingent Allegiance (ECA) is cleared. Processing of queued commands is resumed when a CA/ECA condition is cleared.

(continued)

MODE SELECT (6) Command (15h) (continued)

Table 5-37 Control Mode Page — Field Descriptions (continued)

Field Value	Default	Description
<u>DQue</u>	0	Disable Queue. When set to the default value of 0, the drive supports tagged command queuing. When the Disable Queue bit is set to 1, tagged command queuing is disabled. A subsequent tagged message from the initiator is rejected with a MESSAGE REJECT message and the I/O process is executed as an untagged command.
<u>EECA</u>	0	AEN enabled.
RAC	0	Report A Check Condition bit. Ignored by the drive.
RAERP	0	Ready AER Permission. Not supported by the drive.
UAAERP	0	Unit Attention AER Permission. Not supported by the drive.
EAERP	0	Error AER Permission. Not supported by the drive.
Ready AER Holdoff Period	0	Not supported by the drive.
Busy Timeout Period	FFFFh	Not supported by the drive.

- NOTES:**
1. If the Queue Algorithm Modifier specifies restricted re-ordering (0000b), commands are not allowed to execute concurrently. If this field specifies unrestricted re-ordering (0001b), concurrent I/O execution is allowed.
 2. Fully supported and temporary fields are underlined. Unsupported fields are not underlined. The PS bit is underlined, meaning it will be set on a MODE SENSE command, since the page is savable

MODE SELECT (6) Command (15h) (continued)

5.14.3.8 Notch and Partition Page (0Ch)

The Notch and Partition page contains parameters for drives that implement a variable number of sectors per cylinder, and, support this page. Each section of the logical unit with a different number of sectors is referred to as a *notch* (or band). Each notch must span a set of consecutive logical blocks in the logical unit, the notches must not overlay, and no logical blocks can be excluded from the notches.

Bit Byte	7	6	5	4	3	2	1	0
0	<u>PS</u>	Rsvd	Page Code (0Ch)					
1	Page Length (16h)							
2	ND	<u>LPN</u>	Reserved					
3	Reserved							
4 – 5	Maximum Number of Notches							
6 – 7	<u>Active Notch</u>							
8 – 11	Starting Boundary							
12 – 15	Ending Boundary							
16 – 23	Pages Notched							

Figure 5-34 Notch and Partition Page — Data Format

MODE SELECT (6) Command (15h) (continued)

Table 5-38 Notch and Partition Page — Field Descriptions

Data Field	Default Value	Description
<u>PS</u>	0	Parameters Savable. This bit is only used with the MODE SENSE command. The returned value of 0 indicates that the target is not capable of saving this page in a non-volatile vendor specific location. This bit is reserved with the MODE SELECT command.
ND	1	Notched (banded) Drive. A value of 0 indicates that the drive is not notched and that all other parameters in this page are returned as 0. The default value of 1 indicates that the device is notched. For each supported active notch value, this page defines the starting and ending boundaries of the notch.
<u>LPN</u>	0	Logical or Physical Notch. The default value of 0 indicates that the boundaries of the notch are based on the physical parameters of the logical unit. The cylinder is considered most significant and the head least significant. A value of 1 indicates that the notch boundaries are based on logical blocks of the logical unit.
MaxNotch	24	This field indicates the maximum number of notches supported by the logical unit. The value cannot be changed.
<u>Active Notch</u>	0	The Active Notch field indicates the notch referenced by this, and subsequent MODE SELECT and MODE SENSE commands. The value in this field can be changed by a later MODE SELECT command. The value of the active notch must be = 0 and = 24. An active notch of 0 indicates that this and subsequent MODE SELECT and MODE SENSE commands refer to the parameters that apply across all notches.
Starting Boundary	000000h, 00h	This field indicates the beginning of the active notch or, if the active notch is 0, the beginning boundary of the logical unit. If the LPN value is 1, the starting boundary is a Logical Block Address. If the LPN value is 0, the three most significant bytes designate the starting cylinder number and the least significant byte is the starting head number. This field is unchangeable. When used with the MODE SELECT command, this field is ignored.

NOTE: Temporary fields are underlined. Unsupported fields are not underlined.

(continued)

MODE SELECT (6) Command (15h) (continued)

Table 5-38 Notch and Partition Page — Field Descriptions (continued)

Data Field	Default	
	Value	Description
Ending Boundary	9GB: 0043B9h, 02h	The ending boundary field either indicates the ending of the active notch or, if the active notch is 0, it indicates the ending of the logical unit. If the LPN value is 1, the ending boundary is a Logical Block Address. If the LPN value is 0, the three most significant bytes identify the ending cylinder number and the least significant byte identify the ending head number. This field is unchangeable. When used with the MODE SELECT command, this field is ignored.
	18GB: 0043B9h, 04h	
	36GB: 0043B9h, 09h	
	72GB: 0043B9h, 13h	
Pages Notched	00000000b, 00000000b, 00000000b, 00000000b, 00000000b, 00000000b, 00010000b, 00001000b	<p>This is a 64-bit bit map that indicates which MODE command pages contain parameters that may be different for different notches. The most significant bit (MSB) corresponds to Page 3Fh (Vendor-Unique Caching Page), and the least significant bit (LSB) corresponds to Page 00h (Unit Attention Control Page).</p> <p>If a bit is 1, the corresponding mode page contains parameters that may be different for different notches. If a bit is 0, the corresponding mode page contains parameters that are constant for all notches.</p> <p>The value of 1 in bits 3 and 12 indicates that Page 03h and 0Ch contain different parameters for different notches.</p>

MODE SELECT (6) Command (15h) (continued)

5.14.3.9 XOR Control Mode Page (10h)

The XOR Control Mode Page provides the initiator with the means to obtain/modify certain XOR operating parameters of the target.

Bit	7	6	5	4	3	2	1	0
Byte								
0	<u>PS</u>	Rsvd	Page Code (10h)					
1	Page Length (16h)							
2	Reserved					<u>XOR</u>	<u>Dis</u>	Rsv'd
3	Reserved							
4 - 7	Maximum XOR Write Size							
8 - 11	(MSB)	Reserved						(LSB)
12 - 15	(MSB)	Maximum REGEN Size						(LSB)
16 - 19	(MSB)	Maximum REBUILD READ Size						(LSB)
20 - 21	(MSB)	Reserved						(LSB)
22 - 23	(MSB)	REBUILD Delay						(LSB)

NOTE: Fully supported fields are underlined. Ignored fields are not underlined.

Figure 5-35 XOR Control Mode Page - Data Format

MODE SELECT (6) Command (15h) (continued)

Table 5-39 XOR Control Page - Field Descriptions

Field	Default Value	Description
<u>PS</u>	1	Parameters Savable. This bit is only used with the MODE SENSE command. A returned value of 1 indicates that the target is capable of saving the page in a non-volatile, vendor-specific location. The bit is reserved with the MODE SELECT command
<u>XOR Dis</u>	0	XOR Disable. An XOR disable bit of 0 enables the XOR functions within a device. An XOR bit of 1 disables the XOR functions within a device.
Maximum XOR Write Size	512	This field specifies the maximum transfer length in sectors that the target accepts for the XDWRITE or XPWRITE commands.
Maximum REGEN Size	0	The drive does not support this option and ignores any value placed in this field.
Maximum REBUILD Size	0	The drive does not support this option and ignores any value placed in this field.
REBUILD Delay	0	The drive does not support this option and ignores any value placed in this field.

NOTE: Fully supported fields are underlined. Ignored fields are not underlined.

MODE SELECT (6) Command (15h) (continued)

5.14.3.10 Power Condition Page (1Ah)

The Power Condition Page controls disk drive Power Management functions. Power Management provides three sub-states of UNIT READY:

- **ACTIVE** This is the state of highest power consumption for the disk drive. A media access command is acted upon immediately by the drive.
- **IDLE** This is a state of lower power consumption for the disk drive. A media access command is acted upon after a delay of approximately one second.
- **STANDBY** This is the state of lowest power consumption for the disk drive. A media access command is delayed by approximately 10 seconds before being acted upon.

Note that in all three sub-states, the drive remains READY.

The figure that follows shows the data format of the page; any fully supported fields are underlined meaning that they are savable if modified with a MODE SELECT command with the PS bit set. Following the figure, a table is provided that contains a description of the Power Condition Page fields.

MODE SELECT (6) Command (15h) (continued)

Bit	7	6	5	4	3	2	1	0
Byte								
0	PS	Page Code (1Ah)						
1	Page Length (0Ah)							
2	Reserved							
3	Reserved					IdlEnb	SbyEnb	
4 - 7	Idle Timer							
8 - 11	Standby Timer							

Figure 5-36 Power Condition Page - Data Format

MODE SELECT (6) Command (15h) (continued)

Table 5-40 Power Condition Page – Field Descriptions

Field	Default	
	Value	Description
PS	0	Parameters Savable. This bit is only used with the MODE SENSE command. A returned value of 1 indicates that the target is capable of saving the page in a non-volatile, vendor-specific location. The bit is reserved with the MODE SELECT command.
IdlEnb	0	Enable Idle Timer. This is a lower power consumption state of the drive. A media access command waits for an approximately one second transition time delay from the IDLE to the ACTIVE state.
SbyEnb	0	Enable Standby Timer. Any media access command waits for an approximately 10 second transition time delay from the STANDBY to the ACTIVE state.
Idle Timer	0	The Idle Timer specifies the time, in 100-millisecond increments, before entering the IDLE state after the last media access command. The IdlEnb field must be set to 1. The internal minimum value for this field is 60 seconds.
Standby Timer	0	The Standby Timer specifies the time, in 100-millisecond increments, before entering the Standby state after the last media access command. The SbyEnb field must be set to 1. The internal minimum value for this field is 15 minutes.

NOTE: Fully supported and temporary fields are underlined. Unsupported fields are not underlined.

MODE SELECT (6) Command (15h) (continued)

5.14.3.11 Information Exceptions Control Page (1Ch)

This page defines the methods used by the disk drive to control the reporting and operations of specific informational exception conditions. This page applies only to informational exceptions that report an additional sense code of FAILURE PREDICTION THRESHOLD EXCEEDED to the application client.

Informational exception conditions occur as the result of vendor-specific events within a disk drive. An informational exception condition may occur asynchronous to any commands issued by an application client.

Bit	7	6	5	4	3	2	1	0
Byte								
0	<u>PS</u>	Page Code (1Ch)						
1	Page Length (0Ah)							
2	<u>Perf</u>	Reserved	<u>EWASC</u>	<u>DExcpt</u>	<u>Test</u>	Rsv'd	<u>LogErr</u>	
3	Reserved				<u>MRIE</u>			
4 - 7	(MSB)	<u>Interval Timer</u>						(LSB)
8 - 11	(MSB)	<u>Report Count</u>						(LSB)

Figure 5-37 Information Exceptions Control Page — Data Format

Table 5-41 Information Exceptions Control Page — Field Descriptions

Field	Default Value	Description
<u>PS</u>	1	Parameters Savable. This bit is only used with the MODE SENSE command. A returned value of 1 indicates that the target is capable of saving the page in a non-volatile, vendor-specific location. This bit is reserved with the MODE SELECT command.
<u>PERF</u>	1	Performance. This bit enables/disables performance-sensitive functions. For Atlas 10K II, this means enabling or disabling flushing of SMART data to the media (done hourly). The prediction algorithm is designed to reset the current accumulated data approximately every 100,000 blocks read or written per head. When this bit = 0, current accumulated SMART data is preserved across power cycles. When this bit = 1, current accumulated SMART data is lost across power cycles.

(continued)

MODE SELECT (6) Command (15h) (continued)

Table 5-41 Information Exceptions Control Page — Field Descriptions (continued)

Field	Default Value	Description
<u>EWASC</u>	1	Enable Warning ASC. When this bit = 0, reporting of thermal warning is disabled. When this bit = 1, thermal warning reports are enabled.
<u>DExcpt</u>	0	Disable Exceptions. This bit is used to enable or disable SMART exception reporting. When this bit = 0, SMART exception reporting is enabled. When this bit = 1, all SMART functions are disabled, including reporting SMART exceptions, running the prediction algorithm, and accumulating SMART data.
<u>Test</u>	0	When this bit = 0, the test function is disabled. When this bit = 1, the disk drive reports false SMART trip reports as prescribed in the MRIE field (see below). Interval timer field and report count and method fields apply. When the test bit forces a SMART failure, the fails reported with a sense code/qualifier of 5D/FF.
<u>LogErr</u>	0	When this bit = 0, it indicates that log sense information is vendor-unique. NOTE: LogErr = 0 is the only value supported.
<u>MRIE</u>	4	Method of Reporting Informational Exceptions. The value within this field indicates the method that will be used by the disk drive to report informational exception conditions, as shown in Table 5-41. The relative priority of reporting of multiple informational exceptions is vendor-specific.

(continued)

MODE SELECT (6) Command (15h) (continued)

Table 5-41 Information Exceptions Control Page — Field Descriptions (continued)

Field	Default Value	Description
<u>Interval Timer</u>	0	<p>This field indicates the period, in 100 millisecond increments, for reporting that an informational exception condition has occurred. The disk drive will not report such conditions more frequently than the time specified by the Interval Timer field and as soon as possible after the timer interval has elapsed. After the condition has been reported, the interval timer restarts. A value of 0 in the Interval Timer field indicates that the disk drive will only report the informational exception condition once. A value of FFFFFFFh in the Interval Timer field indicates that the timer interval is vendor-specific.</p> <p>If the disk drive's predictor detects a SMART threshold trip, it reports the failure only if it has not reported a failure within the last interval timer period. A value of 0 means that there is no required waiting between SMART reports.</p> <p>NOTE: The Interval Timer field does not apply to MRIE mode 6.</p>
<u>Report Count</u>	1	<p>The value contained in this field indicates the number of times to report an informational exception condition to the application client. A value of 0 in this field indicates there is no limit on the number of times the disk drive will report an informational exception condition.</p> <p>Each time the disk drive detects a SMART failure, the SMART status is reported the number of times indicated by the value of the Report Count field. The disk drive then resets its cumulative values and continues with its predicting activity. The SMART threshold trip remains permanent (or until a LOG SELECT command clears it) and can be retrieved at any time using MRIE mode 6.</p> <p>NOTE: The Report Count field does not apply to MRIE mode 6.</p>

MODE SELECT (6) Command (15h) (continued)

Table 5-42 Codes Used by the MRIE Field

Field	Description
0h	No reporting of informational exception condition. This method instructs the disk drive not to report these conditions.
1h	Asynchronous Event Reporting. This method of reporting instructs the disk drive to report informational exception conditions by using the rules for asynchronous event reporting as described in the SCSI-3 Architecture Model and the relevant Protocol Standard. The sense key will be set to RECOVERED ERROR and the additional sense code will indicate the cause of the informational exception condition.
2h	Generate Unit Attention. This method of reporting instructs the disk drive to report informational exception conditions by returning a CHECK CONDITION status on any command. The sense key is set to UNIT ATTENTION and the additional sense code will indicate the cause of the information exception condition. The command that has the CHECK CONDITION will not be executed before the informational exception condition is reported.
3h	Conditionally Generate Recovered Error. This method of reporting instructs the disk drive to report informational exception conditions dependent on the value of the PER bit of the Error Recovery Parameters Mode Page, by returning a CHECK CONDITION status on any command. The sense key will be set to RECOVERED ERROR and the additional sense code will indicate the cause of the informational exception condition. The command that has the CHECK CONDITION will complete before and informational exception condition is reported.
4h	Unconditionally Generate Recovered Error. This method of reporting instructs the disk drive to report informational exception conditions regardless of the value of the PER bit of the Error Recovery Parameters Mode Page, by returning a CHECK CONDITION status on any command. The sense key will be set to RECOVERED ERROR and the additional sense code will indicate the cause of the informational exception condition. The command that has the CHECK CONDITION will complete without error before any informational exception condition will be reported.
5h	Generate No Sense. This method of reporting instructs the disk drive to report informational exception conditions by returning a CHECK CONDITION status on any command. The sense key will be set to NO SENSE and the additional sense code will indicate the cause of the informational exception condition. The command that has the CHECK CONDITION will complete without error before any informational exception condition is reported.
6h	Only Report Informational Exception Condition On Request. This method of reporting instructs the disk drive to preserve the informational exception(s) data. To learn about these exception conditions, the application client polls the disk drive via an unsolicited REQUEST SENSE command. The sense key is set to NO SENSE and the additional sense code will indicate the cause of the informational exception condition.
7h – Bh	Reserved.
Ch – Fh	Vendor-specific.

MODE SELECT (6) Command (15h) (continued)

5.14.3.12 Quantum-Unique Page (39h)

This Quantum-unique page is implemented for historical purposes only. All unused fields are changeable but are ignored by the drive.

Bit	7	6	5	4	3	2	1	0
Byte								
0	<u>PS</u>	Rsv'd	Page Code (39h)					
1	Page Length (14h)							
2	<u>Wait SpinUp Enable</u>	<u>Enable Error Logging</u>	<u>Disable Parity Check</u>	<u>Reall. Uncor. Errors</u>	<u>Enable Fill Data Pattern</u>	<u>Disable SCAM</u>	<u>Disable Unit Att'n</u>	<u>Disable IWR</u>
3	<u>SSSIE</u>	<u>Disable Auto WRITE</u>	<u>Disable Auto READ</u>	<u>Delay Spin</u>	<u>Software Selectable SCSI ID</u>			
4	<u>Reselection Retries</u>							
5	<u>ERSS</u>	<u>Stagger Spin</u>	<u>Auto Master Enable</u>	<u>Force 8</u>	<u>DFCP</u>	<u>PXM</u>	<u>ISN</u>	<u>IWN</u>
6 - 7	(MSB) <u>Stagger Spin Delay</u> (LSB)							
8	<u>ASDP</u>	<u>Enable Quiet</u>	Rsv'd	<u>USCIO</u>	<u>SAE</u>	<u>DRC</u>	<u>NMAIOP</u>	Rsv'd
9	<u>Soft Fault</u>	<u>ManCal</u>	<u>DPER</u>	<u>FAxRE</u>	<u>QDER</u>	Rsv'd	<u>DFC</u>	<u>Read Ahead Disable</u>
10	<u>EDAF</u>	<u>Disable Revector Cache</u>	<u>ESPDIs</u>	<u>DFRCSU</u>	<u>AIP Gen Disabled</u>	<u>AIP Check Enabled</u>	<u>AIP Capable</u>	<u>LogPgB</u>
11	<u>ESL</u>	<u>IMR</u>	<u>IOE</u>	<u>ArcsOn</u>	<u>DlyESP</u>	<u>Force 10</u>	<u>DEWG</u>	<u>Enable Comm Timeout</u>
12	<u>MaxSkip</u>							
13 - 14	<u>Vendor Specific</u>							
15	<u>Minimum Block Threshold</u>							
16	<u>Maximum Slots to be Reserved</u>						<u>SSSO</u>	<u>EAMNR</u>
17	<u>Spindle Synch Check Interval</u>							
18 - 19	Reserved							
20	<u>Command Timeout Limit</u>							
21	<u>Command Timeout Limit Response Time</u>							

Figure 5-38 Quantum-Unique Page – Data Format

MODE SELECT (6) Command (15h) (continued)

Table 5-43 Quantum-Unique Page — Field Descriptions

Field	Default	Description
<u>PS</u>	1	Parameters Savable. This bit is only used with the MODE SENSE command. A returned value of 1 indicates that the disk drive is capable of saving the page in a non-volatile, vendor-specific location. The bit is RESERVED with a MODE SELECT command.
<u>Wait SpinUp Enable</u>	0	When this field = 0, the disk drive spins up upon power on. When this field = 1, the disk drive performs a Wait/Spin.
<u>Enable Error Logging</u>	0	When this field = 0, Error Logging is disabled. When this field = 1, Error Logging is enabled.
<u>Disable Parity Check</u>	0	When this field = 0, Parity Checking is enabled. When this field = 1, Parity Checking is disabled.
<u>Reall. Uncorr. Errors</u>	0	When this field = 0, reallocation of blocks that have uncorrectable READ errors does not take place. (Note that if this field is set to 0 and ARRE bit in Page 01 of the Error Recovery Parameters is set to 1, no reallocation takes place.) When this field = 1, and ARRE bit in Page 01 of the Error Recovery Parameters is set to 1, the disk reallocates a block that has an uncorrectable READ error the next time this block is written (the default for ARRE is 1).
<u>Enable Fill Data Pattern</u>	1	When this field = 0, the disk drive ignores any information in Byte 2 of the FORMAT UNIT command. Note that if the setting in Byte 2 of FORMAT UNIT is "0", the disk drive will interpret that setting as "1." When this field = 1, the disk drive writes the data pattern that is specified in Byte 2 of the FORMAT UNIT command into every user-accessible block when the disk drive executes a FORMAT UNIT command.
<u>Disable SCAM</u>	1 (Parallel)	When this field = 0, SCAM is enabled. When this field = 1, SCAM features are disabled.
<u>Disable Unit Att'n</u>	0	When this field = 0, a UNIT ATTENTION condition exists following power on or a reset caused by either a BUS DEVICE RESET message or a hard reset condition. When this field = 1, the disk drive will not issue a CHECK CONDITION status with the UNIT ATTENTION sense key on the first command following a reset. The disk drive will execute the first command issued after a reset condition.
<u>Disable IWR</u>	0	Disable Wide Residue Message. When this field = 0, the disk drive sends the IGNORE WIDE RESIDUE message. When this field = 1, the disk drive will not send the IGNORE WIDE RESIDUE message.

(continued)

MODE SELECT (6) Command (15h) (continued)

Table 5-43 Quantum-Unique Page — Field Descriptions (continued)

Field	Default	Description
<u>SSSIE</u>	0	Software Selectable SCSI ID Enable. When this field = 0, the disk drive's SCSI ID is determined via the disk drive's hardware ID select jumper settings. When this field = 1, the disk drive's SCSI ID setting is determined by the value set in the Software Selectable SCSI ID field (Byte 3, bits 0 – 3).
<u>Disable Auto WRITE</u>	0	When this field = 0, the Auto WRITE function is enabled. When this field = 1, the Auto WRITE function is disabled.
<u>Disable Auto READ</u>	0	When this field = 0, the Auto WRITE function is enabled. When this field = 1, the Auto WRITE function is disabled.
<u>Delay Spin</u>	0	When this field = 0, Delay Spin is not enabled. When this field = 1, and Staggered Spin Delay (Bytes 6 & 7) is enabled, the three lower bits of the disk drive's SCSI ID are used for the multiplier to determine the length of the delay. When only three bits are used, disk drives 0 through 7 and drives 8 through 15 have the same corresponding delays since the fourth bit is not used (disk drive 0 has the same delay as disk drive 8, disk drive 1 has the same delay as disk drive 9, and so on).
<u>Software Selectable SCSI ID</u>	0	When SSSIE (Byte 3, bit 7) = 1, this field determines the disk drive's SCSI ID. When changing SCSI ID via the MODE SELECT command, the SCSI ID changes after the next reset or power cycle. If at initial power on the SSSIE bit is set to 1 (the bit's previous setting was saved to disk) the Software Selectable SCSI ID field's value determines the disk drive's SCSI ID.
<u>Reselection Retries</u>	7	The value of this field specifies the number of reselection attempts the drive will make to the initiator. A value of 0 indicates an infinite number of reselection attempts.
ERSS	0	Enable Rotational Position Locking (RPL) Sector Synchronization. Not used by the drive.
<u>Stagger Spin</u>	0	When the value of this field = 1, the Stagger Spin Delay field (Bytes 6 & 7) is in effect. The disk drive operates as if its Spin Delay jumper was installed.

(continued)

MODE SELECT (6) Command (15h) (continued)

Table 5-43 Quantum-Unique Page — Field Descriptions (continued)

Field	Default	Description
<u>Auto Master Enable</u>	0	When this field = 0, this disk drive does not act as the master in a storage array configuration (refer to the Hard Disk Geometry page). When this field = 1, this disk drive becomes the master. Not used by the drive.
<u>Force 8</u>	0	When this field = 0, 16-bit wide data transfer responses are allowed. When this field = 1, negotiations are limited to 8-bit responses.
<u>DFCP</u>	0	Delay Following Command Phase. Not used by the drive.
<u>PXM</u>	0	Preserve Synchronous and Wide Mode. If set = 1, this bit causes the disk drive to preserve the most recent wide and synchronous states for all initiators across power on resets, bus reset, and bus device reset messages. Not used by the drive.
<u>ISN</u>	0	Initiate Synchronous Negotiation. When this field is set = 1, the disk drive initiates a synchronous transfer negotiation with the initiator. This is done on the first selection with ATTENTION following a power on reset, bus reset, or bus device reset. The negotiation is done for each initiator. Note that if both ISN and IWN (Byte 5, bit 0) are set = 1, the wide data transfer negotiation takes place first.
<u>IWN</u>	0	Initiate Wide Negotiation. When this field is set = 1, the disk drive initiates a wide data transfer negotiation with the initiator. This is done on the first selection with ATTENTION following a power on reset, bus reset, or bus device reset. The negotiation is done for each initiator. Note that if the disk drive is not capable of actually transferring data on a wide bus, the firmware will only negotiate for an 8-bit bus. If both ISN (Byte 5, bit 1) and IWN are set = 1, the wide data transfer negotiation takes place first. A wide data transfer negotiation causes the disk drive to reset any synchronous agreement to asynchronous mode.

(continued)

MODE SELECT (6) Command (15h) (continued)

Table 5-43 Quantum-Unique Page — Field Descriptions (continued)

Field	Default	Description
<u>Stagger Spin Delay</u>	28h	<p>The Stagger Spin Delay is a value in 100 millisecond units that is used to determine the length of the period of time delay before starting the disk drive's motor if the Stagger Spin feature is used. The actual amount of the delay is determined using the following formula:</p> $\text{Spin Up Delay} = (\text{SCSI ID} \times \text{Stagger Spin Delay}).$ <p>Refer to the description of Delay Spin (Byte 3, bit 4).</p> <p>Note that for the value in Stagger Spin Delay field to have effect, the Spin Delay hardware jumper must be installed (Chapter 3).</p>
<u>ASDP</u>	0	Always Save Data Pointers. If set = 1, the drive will always send an SDP message before sending a DISCONNECT message.
<u>Enable Quiet</u>	0	Enable Quiet Seek Mode. If set = 1, Quiet Seek mode is enabled.
USCIO	0	Use Same Cylinder Optimization in OPTWORK. When set = 1, helps multiple sequential streams.
<u>SAE</u>	0	Show All Errors. When this field = 1, errors usually hidden are shown.
<u>DRC</u>	0	<p>Disable Reassign Copy. When this field = 0, the disk drive copies best-guess data to the new physical block during REASSIGN BLOCKS command processing. If the disk drive cannot successfully read the original block's data, the disk drive sets a Forced Error flag in the replacement block's flags.</p> <p>When this field = 1, the disk drive does not copy a block's data and obtain best-guess data; it first reassigns the block and then writes data to the block. Until it is overwritten by a subsequent WRITE command, a reassigned block's data is undefined and its Forced Error flag is set.</p>
<u>NMAIOP</u>	0	No MAIOP. When this field = 1, MAIOP is disabled.
<u>DNTRA</u>	0	Disable Next Track Read Ahead. Do not prefetch next track.
<u>Soft Fault</u>	0	When this field is set = 1, the disk drive's FAULT led illuminates.

(continued)

MODE SELECT (6) Command (15h) (continued)

Table 5-43 Quantum-Unique Page — Field Descriptions (continued)

Field	Default	Description
<u>ManCal</u>	0	Manual Calibration. When set = 0, the disk drive runs calibration and diagnostic procedures periodically. When set = 1, calibrations occur only as a result of a REZERO UNIT command. Note that this bit is not saved by MODE SELECT when the value of the SP field is 1.
<u>DPER</u>	1	Disable Post Error Recovery. When set = 1, disables Post Error Recovery (refer to Read-Write Error Recovery Page, Section 5.17.3.3).
<u>FAXRE</u>	1	Force AxRE where x = W (WRITE) or R (READ) When set = 0, the settings of AWRE/ARRE are as set in Mode Page 01h (Read-Write Error Recovery Page). When set = 1, Automatic WRITES or READs are forced, even if the related fields of Mode Page 01h are = 0.
<u>QDER</u>	0	Quick Disk Error Recovery. When set = 0, disk error recovery has no 20msec limit. When set = 1, disk error recovery time allotted is less than 20msec.
<u>DFC</u>	0	Disable Format Certification. If set = 1, format certification is disabled regardless of the value of the DCRT bit in the FORMAT Command's parameter header.
<u>Read Ahead Disable</u>	0	If set = 1, disables READ ahead.
<u>Enable Double Ack</u>	0	If set = 1, disables the double acknowledge filter.
<u>ESPDIs</u>	0	ESP Disabled. If this field = 1, ESP (anticipatory READ) is disabled.
<u>DFRCSU</u>	0	Disable Filling of Revector Cache on Spin-Up. When this field = 1, revectoring cache is not refilled on spin-up.
<u>AIP Gen Disable</u>	0	Asynchronous Information Protection (AIP) Generation Disabled. When set = 1, AIP generation is disabled.
<u>AIP Check Enabled</u>	0	AIP Checking Enabled. The drive does not support AIP Checking.
AIP Capable	1	If = 1, indicates drive is capable of AIP.
<u>LogPgB</u>	0	When this field = 0, "total bytes read/written" field on Log Page 02h and 03h reported as number of <i>bytes</i> . When = 1, field is reported as number of <i>blocks</i> .

(continued)

MODE SELECT (6) Command (15h) (continued)

Table 5-43 Quantum-Unique Page — Field Descriptions (continued)

Field	Default	Description
<u>ESL</u>	0	Enable Save Log. When set = 1, saving and loading of non-SMART pages to/from the MCS area is enabled.
<u>IMR</u>	0	Infinite Message Retry. When set = 1, the number of retries at a message error is infinite.
<u>IOE</u>	1	Individual ORCA Profile Enabled. When set = 0, this feature is disabled.
<u>ArcsOn</u>	0	Arcs always on.
<u>DlyESP</u>	0	Delay ESP. For ESP (anticipatory-READ) commands only, when set = 1 the delay from data phase to the first REQ of data is increased from about 2 μ s to about 16 μ s.
<u>Force10</u>	0	When this bit = 0, the disk drive may negotiate for a transfer period less than 19h. When this bit = 1, SDTR transfer period is limited to 19h minimum.
<u>DEWG</u>	1	Disable ESP WRITE Guess. When = 1, enables this feature.
<u>Enable Comm Timeout</u>	0	Enable Command Timeout. When = 1, enables this feature.
<u>MaxSkip</u>	5	The value of this field specifies the number of times that a host request can be skipped over (such as when performing queuing optimizations). This parameter affects only queued commands. Some host operating systems support tagged command queuing which allows for multiple requests to be resident in the disk drive simultaneously. The disk drive can then re-order execution of the requests to improve throughput. For systems in which minimum completion time of a request is more important than throughput, MaxSkip can be reduced. Values allowed range from 0 to 255.
<u>Vendor Specific</u>	---	Proprietary Information
<u>Minimum Block Threshold</u>	1	This field specifies the minimum number of blocks ready before starting data transfer.

(continued)

MODE SELECT (6) Command (15h) (continued)

Table 5-43 Quantum-Unique Page — Field Descriptions (continued)

Field	Default	Description
<u>Maximum Slots to be Reserved</u>	0	This field is used to reserve x amount of slots from the total resources within the disk drive for incoming commands. Once the remaining slots are full, up to the queue depth of the disk drive, subsequent commands will be accepted but placed on a disk drive internal pending list. As the disk drive completes the active commands, freeing up internal resources (SMBs), the pending commands are activated.
<u>SSSO</u>	0	Spindle Sync Sector Offset. This field is used to specify the RPL offset to sector boundary. This drive does not support spindle synchronization.
<u>EAMNR</u>	0	Enable Auto Master Negotiation if RPL Mode = 0. This drive does not support spindle synchronization.
Spindle Sync Check Interval	0	The value of this field specifies the time, in minutes, of the interval between spindle synchronization status checks. If this value is 0, the spindle is checked at five-minute intervals. If the value is set to 60 or greater, the spindle is checked at 60-minute intervals. This drive does not support spindle synchronization.
<u>Command Timeout Limit</u>	60	The value of this field specifies the timeout limit in 100 millisecond units. The default value of 60 specifies a command timeout limit of 6 seconds
<u>Spindle Sync Check Interval</u>	10	The value of this field specifies how long, in 100 millisecond units, before the Command Timeout Limit abort processing should commence. This is intended to compensate for abort processing time.

NOTE: Fully supported and temporary fields are underlined. Ignored fields are not underlined. The PS bit is underlined, meaning it will be set on a MODE SENSE command, since the page is savable.

Asynchronous Information Protection (AIP)

To start AIP:

1. Make sure "AIP Generation Disabled" (byte 10, bit 3) is **cleared**. Default is 0.
2. Make sure "AIP Checking Enabled" (byte 10, bit 2) is set for EACH. This bit is maintained uniquely for each unique initiator ID that senses it or selects it, so all initiators that want to participate in AIP must set this for themselves. Default is 0.
3. Send a negotiation message (SDTR, WDTR, or PPR) with valid AIP data. The drive will check to see if it received valid AIP data from the host for the entire negotiation message. If so, AIP is activated and remains so for as long as that initiator is still valid. If not, no errors are reported, and another attempt is made at the next negotiation.

5.15 MODE SELECT (10) Command (55h)

The MODE SELECT (10) command allows the initiator to modify certain mode pages and thereby control some of the drive's operational characteristics. Refer to MODE SELECT (6) for additional descriptions. See Figure 5-39 for the data format of the MODE SELECT (10) Command Descriptor Block. See MODE SELECT (6), Table 5-44, for descriptions of the fields. See Figure 5-40 for the Mode Parameter Header (10 Byte) — Data Format and Figure 5-44 for the Mode Parameter Block Descriptor — Data Format. See Table 5-25 for mode parameter block field descriptions.

Bit	7	6	5	4	3	2	1	0
Byte								
0	Operation Code (55h)							
1	Reserved			PF	Reserved			SP
2 - 6	Reserved							
7 - 8	Parameter List Length							
9	Control							

Figure 5-39 MODE SELECT (10) Command Descriptor Block – Data Format

Bit	7	6	5	4	3	2	1	0
Byte								
0 - 1	(MSB) Mode Data Length							(LSB)
2	Medium Type							
3	Device-Specific Parameter							
4 - 5	Reserved							
6 - 7	(MSB) Block Descriptor Length							(LSB)

Figure 5-40 Mode Parameter Header (10-Byte) — Data Format

MODE SELECT (10) Command (55h) (continued)

Bit Byte	7	6	5	4	3	2	1	0
0 - 3	(MSB) Number of Blocks							(LSB)
4	Density Code							
5 - 7	(MSB) Block Length							(LSB)

Figure 5-41 Mode Parameter Block Descriptor — Data Format

Table 5-44 Mode Parameter Block Descriptor — Field Descriptions

Data Field	Description
Number of Blocks	The value in this field specifies the number of logical blocks on the disk to which the density code and block length fields apply. A value of 0 in this field indicates that all of the remaining logical blocks on the disk have the medium characteristics specified.
Density Code	The Density Code field is unique for each device type. Direct access devices value is zero.
Block Length	The value in this field specifies the length in bytes of each logical block described by the block descriptor.

This page intentionally left blank.

This page intentionally left blank.

5.16 MODE SENSE (6) Command (1Ah)

The MODE SENSE (6) command allows the initiator to read the drive's sector descriptor and mode pages. The MODE SENSE Command Descriptor Sector is shown in Figure 5-42 and the data fields are described in Table 5-45. The Mode Parameter Header is shown in Figure 5-43 and the data fields are described in Table 5-46.

Bit Byte	7	6	5	4	3	2	1	0
0	Operation Code (1Ah)							
1	Logical Unit Number			Reserved	DBD	Reserved		
2	PC		Page Code					
3	Reserved							
4	Allocation Length							
5	Control							

Figure 5-42 MODE SENSE (6) Command Descriptor Block — Data Format

Table 5-45 MODE SENSE Command — Field Descriptions

Data Field	Description
DBD	Disable Block Descriptor. When the value is 0, the drive returns its Block Descriptor followed by the requested page(s). When the value is 1, the drive returns only the requested pages.
PC	Page Control. Determines which type of page values are returned. These values are: 00b Current 01b Changeable 10b Default 11b Saved
Page Code	The Page Code field determines the pages to be returned to the initiator in the command's data-in buffer transfer. Refer to Tables 5-26 and 5-47 for a complete list of the Mode Pages (changeable and read-only). A Page Code of 3Fh causes all supported pages to be returned.

MODE SENSE (6) Command (1Ah) (continued)

Bit Byte	7	6	5	4	3	2	1	0
0	Mode Data Length							
1	Medium Type							
2	Device-Specific Parameter							
3	Block Descriptor Length							

*Figure 5-43 Mode Parameter Header (6 Byte) — Data Format**Table 5-46 Mode Parameter Header — Field Descriptions*

Data Field	Description
Mode Data Length	Length of the returned data
Medium Type	Always zero (0)
Device-Specific Parameter	Bit 7 - This bit is set if the drive is write protected Bit 4 - This bit is always set to indicate that the drive supports cache memory.
Block Descriptor	Zero (0) if Block Descriptor (DBD) in MODE SENSE Command Descriptor Length sector is 0. If DBD is 1, then the Block Descriptor Length is eight (8).

5.16.1 Read-Only Mode Pages

Table 5-46 lists read-only mode pages. Table 5-25 lists initiator-changeable mode pages.

Table 5-47 Read-Only Mode Pages

Page Code	Page Name	Reports the drive's	Size (bytes)
03h	Format Device	Medium format	24
04h	Rigid Disk Geometry	Total cylinder head and rotational speed	24
38h	Vendor Unique Caching	Total segments	16

*MODE SENSE (6) Command (1Ah) (continued)***5.16.2 Format Device Page (03h)**

The Format Drive Page (Figure 5-44) describes the drive's medium format. The data fields are described in Table 5-48. Note that the page is *savable* but no parameters are *changeable*.

Bit Byte	7	6	5	4	3	2	1	0
0	<u>PS</u>	Rsvd	Page Code (03h)					
1	Page Length (16h)							
2 - 3	Tracks Per Zone							
4 - 5	Alternate Sectors/Zone							
6 - 7	Alternate Tracks/Zone							
8 - 9	Alternate Tracks/Logical Unit							
10 - 11	Sectors/Track							
12 - 13	Data Bytes/Physical Sector							
14 - 15	Interleave							
16 - 17	Track Skew Factor							
18 - 19	Cylinder Skew Factor							
20	SSEC	HSEC	RMB	SURF	Reserved			
21 - 23	Reserved							

Figure 5-44 Format Device Page — Data Format

MODE SENSE (6) Command (1Ah) (continued)

Table 5-48 Format Device Page — Field Descriptions

Field	Value	Description
PS	1	Parameters Savable. This bit is only used with the MODE SENSE command. This bit is reserved with the MODE SELECT command. A PS bit value of 1 indicates that the disk drive is capable of saving the page in a non-volatile vendor specific location.
Tracks/Zone		The number of physical tracks per cylinder: 3 for 9.2 GB model, 5 for 18.4 GB, 10 for 36.7 GB, 20 for 73.4 GB
Alternate Sectors/Zone		The number of spare locations within a notch: 37 for 9.2 GB model, 63 for 18.4 GB, XX for 36.7 GB, XX for 73.4 GB
Alternate Tracks/Zone	0	The drive does not have any alternate tracks.
Alternate Tracks/Logical Unit		The number of tracks reserved for defect handling: 46 for 9.2 GB model, 78 for 18.4 GB, XX for 36.7 GB, XX for 73.4 GB
Sectors/Track	413	The number of Physical Block Address sectors per track. The value depends on the Notch value selected in the Notch and Partitions page.
Data Bytes/Physical Sector	512	This value is the current disk sector size.
Interleave	1	The drive does not support interleave.
Track Skew Factor		The number of physical sectors between the last logical sector on one track and the first logical sector on the next track within a cylinder: 15 for 9.2 GB model, 15 for 18.4 GB, 18 for 36.7 GB, 15 for 73.4 GB
Cylinder Skew Factor		The number of physical sectors between the last logical sector on one cylinder and the first logical sector on the next cylinder: 25 for 9.2 GB model, 20 for 18.4 GB, 20 for 36.7 GB, 25 for 73.4 GB
SSEC	0	The drive does not use soft sector formatting.
HSEC	1	The drive supports hard sector formatting only.
RMB	0	The drive does not support removable logical media.
SURF	0	A Surface bit value of 0 indicates that the drive allocates progressive addresses to all logical blocks within a cylinder prior to allocating addresses on the next cylinder.

*MODE SENSE (6) Command (1Ah) (continued)***5.16.3 Rigid Disk Geometry Page (04h)**

The Rigid Disk Geometry Page describes the drive geometry.

Bit	7	6	5	4	3	2	1	0
Byte								
0	<u>PS</u>	Rsvd	Page Code (04h)					
1	Page Length (16h)							
2 - 4	Number of Cylinders							
5	Number of Heads							
6 - 8	Starting Cylinder - Write Precompensation							
9 - 11	Starting Cylinder - Reduced Write Current							
12 - 13	Drive Step Rate (0)							
14 - 16	Landing Zone Cylinder							
17	Reserved						<u>RPL</u>	
18	<u>Rotational Offset</u>							
19	Reserved							
20 - 21	Medium Rotation Rate							
22 - 23	Reserved							

Figure 5-45 Rigid Disk Geometry Page — Data Format

MODE SENSE (6) Command (1Ah) (continued)

Table 5-49 Rigid Disk Geometry Page — Field Descriptions

Field	Value	Description
<u>PS</u>	1	Parameters Savable. This bit is only used with the MODE SENSE command. A returned value of 1 indicates that the target is capable of saving the page in a non-volatile, vendor-specific location. The bit is reserved with the MODE SELECT command.
Number of Cylinders	17338	The number of physical cylinders that hold host data.
Number of Heads	3/5/10/20	The number of heads that can access host data.
Starting Cylinder - Write Precompensation	17338	The value in this field is the physical cylinder at which write precompensation is to begin. The first cylinder number is 0. If the starting cylinder is equal to the value in the Number of Cylinders field, write precompensation is disabled.
Starting Cylinder - Reduced Write Current	0	The value in this field is the physical cylinder at which write current is reduced. The first cylinder number is 0. If the starting cylinder for reduced write current is equal to the value in the Number of Cylinders field, reduced write current is disabled.
Drive Step Rate	0	Applies to stepper-motor actuators. Not applicable.
Landing Zone Cylinder	17338	This field indicates a two's complement location where the drive positions the disk heads. A negative value indicates that the heads are positioned below cylinder 0 by that number of cylinders. A positive value greater than the number of cylinders indicates that the heads are positioned beyond the cylinders used for data storage at the location specified.
RPL	00b	Rotational Position Locking is used for spindle synchronization in storage arrays. The modes are: 00b – Disabled 10b – Master 01b – Slave 11b – Not supported RPL can be overridden if the Quantum (Vendor) Special Function Control page Auto field value = 1. This drive does not support spindle synchronization.

(continued)

MODE SENSE (6) Command (1Ah) (continued)

Table 5-49 Rigid Disk Geometry Page — Field Descriptions (continued)

Field	Value	Description
Rotational Offset	0	<p>Rotational Offset indicates the amount of rotational offset that the drive uses when synchronized. The rotational offset is applied in the retarded direction (lagging the synchronized spindle master). The value in the field is determined by the following formula:</p> $\text{Offset} = \frac{0 \dots 255}{256}$ <p>For example, a numerator of 128 indicates an offset of 1/2 revolution (180°).</p> <p>The valid parameter values are 0 to 255.</p> <p>The specified value may be rounded. Rotational offset is not used when the drive is configured as the synchronized-spindle master.</p>
Medium Rotation Rate	10021	The disk rotation speed in RPM.

*MODE SENSE (6) Command (1Ah) (continued)***5.16.4 Vendor-Unique Caching Page (38h)**

The Vendor-Unique Caching Page is a read-only page that reports the number of segments in the drive's cache memory.

Bit Byte	7	6	5	4	3	2	1	0
0	PS	Rsv'd	Page Code (38h)					
1	Page Length (0Eh)							
2	Reserved				Cache Table Size			
3 - 15	Reserved							

Figure 5-46 Vendor-Unique Caching Page — Data Format

Table 5-50 Vendor-Unique Caching — Field Description

Field	Value	Description
PS	0	Parameters Savable. This bit is only used with the MODE SENSE command. This bit is reserved with the MODE SELECT command. A PS bit value of 1 indicates that the disk drive is capable of saving the page in a non-volatile vendor specific location.
Cache Table Size	27	This field indicates the number of cache segments supported by the drive.

5.17 MODE SENSE (10) COMMAND (5Ah)

The MODE SENSE command allows the initiator to read the drive's sector descriptor and mode pages.

See Figure 5-47 for MODE SENSE (10) Command Descriptor Block – Data Format, Figure 5-48 for Mode Parameter Header (10 Byte) — Data Format, and Figure 5-49 for Mode Parameter Block Descriptor – Data Format. Refer to Tables 5-45, 5-46, and 5-51 for descriptions of the fields.

Bit Byte	7	6	5	4	3	2	1	0
0	Operation Code (5Ah)							
1	Logical Unit Number			Reserve d	DBD	Reserved		
2 - 6	Reserved							
7 - 8	(MSB) Allocation Length							(LSB)
9	Control							

Figure 5-47 MODE SENSE (10) Command Descriptor Block – Data Format

Bit Byte	7	6	5	4	3	2	1	0
0 - 1	(MSB) Mode Data Length							(LSB)
2	Medium Type							
3	Device-Specific Parameter							
4 - 5	Reserved							
6 - 7	Block Descriptor Length							

Figure 5-48 Mode Parameter Header (10 Byte) — Data Format

Field Descriptions for the Mode Parameter Header are given in the table that follows.

MODE SENSE (10) COMMAND (5Ah) (continued)

Bit Byte	7	6	5	4	3	2	1	0
0	Density Code							
1 - 3	(MSB) Number of Blocks							(LSB)
4	Reserved							
5 - 7	(MSB) Block Length							(LSB)

Figure 5-49 Mode Parameter Block Descriptor — Data Format

Table 5-51 Mode Parameter Header & Block Descriptor — Field Descriptions

Data Field	Description
Mode Data Length	Reserved. Must be 0.
Medium Type	Ignored by the drive.
Device Specific Parameter	Ignored by the drive.
Block Descriptor Length	The value = 0 if no Block Descriptor is supplied; the length is 8 if a block descriptor is supplied.
Density Code	Ignored by the drive.
Number of Blocks	The value in this field specifies the number of logical blocks on the unit that are affected by the current MODE SELECT command. When this field = 0, the MODE SELECT parameters are put into effect for all logical blocks on the unit. The maximum value allowed is the total number of initiator-addressable logical blocks for the value specified in the Block Length field. A value that is less than the maximum value will decrease the number of initiator-accessible logical blocks to the amount specified.
Block Length	The value in this field specifies the length, in bytes, of each logical block.

5.18 PERSISTENT RESERVATION IN Command (5Eh)

The PERSISTENT RESERVATION IN command is a 10-byte command used to obtain information about persistent reservations and reservations that are active within a device server. It is used in conjunction with the PERSISTENT RESERVATION OUT command.

Note that a if device that has RESERVE (10), RESERVE (6), extent reservations, or SCSI Media Changer (SMC) element reservations active receives a PERSISTENT RESERVATION IN command, the PERSISTENT RESERVATION IN command is rejected with a RESERVATION CONFLICT status reported.

The figure below illustrates the format of the PERSISTENT RESERVATION IN command; the table that follows explains the data fields of the command.

Bit	7	6	5	4	3	2	1	0
Byte								
0	Operation Code (5Eh)							
1	Reserved			Service Action				
2 - 6	Reserved							
7 - 8	(MSB) Allocation Length							(LSB)
9	Control							

Figure 5-50 PERSISTENT RESERVATION IN Command Descriptor Block — Data Format

PERSISTENT RESERVATION IN Command (5Eh) (continued)

Table 5-52 PERSISTENT RESERVATION IN Command — Field Descriptions

Data Field	Description												
Service Action	<p>Service actions that require information about persistent reservation and registrations may require enabling of nonvolatile memory within the logical unit.</p> <p>Service action codes available are:</p> <table border="1"> <thead> <tr> <th>Code</th> <th>Name</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00h</td> <td>Read Keys</td> <td>Reads all registered Reservation keys</td> </tr> <tr> <td>01h</td> <td>Read Reservations</td> <td>Reads all current persistent reservations</td> </tr> <tr> <td>02 – 1Fh</td> <td>Reserved</td> <td>Reserved</td> </tr> </tbody> </table> <p>A "Read Keys" service action requests that the device server return a parameter list that includes a header and a complete list of all of the reservation keys currently registered with the device server. If multiple initiators have registered with the same key, then the key is listed multiple times, once for each registration. Refer to Figure 5-51 and Table 5-53 for information about Read Keys parameter data.</p> <p>A "Read Reservation" service action requests that the device server return a parameter list that contains a header and a complete list of all persistent reservations that are presently active in the device server and its extents. Note that duplicate persistent reservations from the same initiator are not reported. Refer to Figure 5-53 and Table 5-54 for information about Read Reservations parameter data.</p>	Code	Name	Description	00h	Read Keys	Reads all registered Reservation keys	01h	Read Reservations	Reads all current persistent reservations	02 – 1Fh	Reserved	Reserved
Code	Name	Description											
00h	Read Keys	Reads all registered Reservation keys											
01h	Read Reservations	Reads all current persistent reservations											
02 – 1Fh	Reserved	Reserved											
Allocation Length	<p>This field indicates how much space has been reserved for the returned parameter list (Read Keys or Read Reservations parameters). The actual length of the parameter data is indicated in the parameter data field for those parameters.</p> <p>If the Allocation Length is not sufficient to contain the entire list of parameters, the first portion of the list that does fit is returned. If it is determined that the remainder of the list is required, the client should send a new PERSISTENT RESERVATION IN command with an Allocation Length field large enough to contain the entire list of parameters.</p>												

PERSISTENT RESERVATION IN Command (5Eh) (continued)

The figure and table below illustrate and describe the data fields of Read Key data parameters.




Figure 5-51 Read Keys Parameters — Data Format

Table 5-53 Read Keys Parameters — Field Descriptions

Data Field	Description
Generation	<p>The value in this field is a 32-bit counter in the device server that is incremented each time a PERSISTENT RESERVATION OUT command requests a Register, Clear, Pre-empt, or Pre-empt and Clear operation. Note that PERSISTENT RESERVATION IN commands do not increment the counter, nor do PERSISTENT RESERVATION OUT commands that perform a Reserve or Release service action, or by a PERSISTENT RESERVATION OUT command that is not done due to an error or a reservation conflict. The value in the Generation field is set to 0 as part of the power on or reset processes.</p> <p>The value in the Generation field allows the application client that examines the value to verify that the configuration of the initiators attached to a logical unit has not been modified by another application client without any notification of the application client doing the examination.</p>
Additional Length	<p>This field contains the count of the number of bytes that are in the Reservation Key list (bytes 8 - n). Note that this field contains the number of bytes in the reservation key list regardless of the value prescribed by the Allocation Length field in the command's CDB.</p>
Reservation Keys	<p>Each of the Reservation Keys appear as items in a list as bytes 8 through n. Each entry reflects an 8-byte reservation key registered with the device server via the PERSISTENT RESERVATION OUT, Reserve, Pre-empt, Pre-empt and Clear, or Register service actions. Each key can be examined by the application client for correlation with a set of initiators and SCSI ports</p>

PERSISTENT RESERVATION IN Command (5Eh) (continued)

The figure and table below illustrate and describe the data fields of Read Reservations data parameters.

Bit	7	6	5	4	3	2	1	0
Byte								
0 - 3	(MSB) Generation (LSB)							
4 - 7	(MSB) Additional Length ($n - 7$) (LSB)							
8 - n	(MSB) Reservation Descriptors (See Figure 5-53 / Table 5-55 for detail of a Descriptor Field) (LSB)							

Figure 5-52 Read Reservations Parameters — Data Format

Table 5-54 Read Reservations Parameters — Field Descriptions

Data Field	Description
Generation	The value in this field is a 32-bit counter in the device server that is incremented each time a PERSISTENT RESERVATION OUT command requests a Register, Clear, Pre-empt, or Pre-empt and Clear operation. Note that PERSISTENT RESERVATION IN commands do not increment the counter, nor do PERSISTENT RESERVATION OUT commands that perform a Reserve or Release service action, or by a PERSISTENT RESERVATION OUT command that is not done due to an error or a reservation conflict. The value in the Generation field is set to 0 as part of the power on or reset processes. The value in the Generation field allows the application client that examines the value to verify that the configuration of the initiators attached to a logical unit has not been modified by another application client without any notification of the application client doing the examination.
Additional Length	This field contains the count of the number of bytes of Reservation descriptors (bytes 8 - n). Note that this field contains the number of bytes regardless of the value prescribed by the Allocation Length field in the command's CDB.
Reservations Descriptors	One Reservation descriptor is reported for each unique persistent reservation on the logical unit when the PERSISTENT RESERVATION IN command has indicated a Read Reservations action. Figure 5-54 and Table 5-55 detail the contents of each Reservation Descriptors field.

PERSISTENT RESERVATION IN Command (5Eh) (continued)

The figure and table below illustrate and describe the data fields of each Read Reservations descriptor's data fields.

Bit Byte	7	6	5	4	3	2	1	0
0 - 7	(MSB) Reservation Key (LSB)							
8 - 11	(MSB) Scope-Specific Address (LSB)							
12	Reserved							
13	Scope				Type			
14 - 15	(MSB) Extent Length (LSB)							

Figure 5-53 PERSISTENT RESERVATION IN Read Reservations Descriptor — Data Format

Table 5-55 PERSISTENT RESERVATION IN Read Reservations Descriptor — Field Descriptions

Data Field	Description
Reservation Key	The Reservation key field contains an 8-byte value that identifies the reservation key under which the persistent reservation is held.
Scope-Specific Address	If the Scope field (Byte 13, bits 4 – 7) represents an “Extent” reservation, the Scope-Specific Address field contains the logical block address (LBA) of the first block of the extent and the Extent Length field (Bytes 14 – 15) contains the number of blocks in the extent. If the Scope field represents an “Element” reservation, the Scope-Specific Address field contains the Element address, zero-filled in the most significant bytes to fit the field and the Extent Length field is set to zero.

(continued)

PERSISTENT RESERVATION IN Command (5Eh) (continued)

Table 5-55 PERSISTENT RESERVATION IN Read Reservations Descriptor — Field Descriptions (continued)

Data Field	Description																	
Scope	<p>The value in this field indicates whether a persistent reservation applies to an entire logical unit, to a part of the logical unit (defined as an extent), or to an element.</p> <p>The values for the Scope field are:</p> <table border="1"> <thead> <tr> <th>Code</th> <th>Name</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0h</td> <td>LU</td> <td>Logical Unit. Persistent reservation applies to the full logical unit. The LU scope is therefore implemented by all device servers that implement PERSISTENT RESERVATION OUT.</td> </tr> <tr> <td>1h</td> <td>Extent</td> <td>Persistent reservation applies to the specified extent. When Extent is the scope, it indicates that the persistent reservation applies to the extent of the logical unit defined by the Scope-Specific Address and Extent Length fields in the PERSISTENT RESERVATION OUT command's parameter list. Note that an extent is defined only for devices defining contiguous logical block addresses. The Extent scope is optional for all device servers that implement PERSISTENT RESERVATION OUT.</td> </tr> <tr> <td>2h</td> <td>Element</td> <td>Persistent reservation applies to the specified element. When Element is the scope, it indicates that the persistent reservation applies to the element of the logical unit defined by the Scope-Specific Address field in the PERSISTENT RESERVATION OUT parameter list. An element is as defined by the SCSI-3 Medium Changer Commands (SMC) standard. The Element scope is optional for all device servers that implement PERSISTENT RESERVATION OUT.</td> </tr> <tr> <td>3h – Fh</td> <td>Reserved</td> <td>Reserved</td> </tr> </tbody> </table>			Code	Name	Description	0h	LU	Logical Unit. Persistent reservation applies to the full logical unit. The LU scope is therefore implemented by all device servers that implement PERSISTENT RESERVATION OUT.	1h	Extent	Persistent reservation applies to the specified extent. When Extent is the scope, it indicates that the persistent reservation applies to the extent of the logical unit defined by the Scope-Specific Address and Extent Length fields in the PERSISTENT RESERVATION OUT command's parameter list. Note that an extent is defined only for devices defining contiguous logical block addresses. The Extent scope is optional for all device servers that implement PERSISTENT RESERVATION OUT.	2h	Element	Persistent reservation applies to the specified element. When Element is the scope, it indicates that the persistent reservation applies to the element of the logical unit defined by the Scope-Specific Address field in the PERSISTENT RESERVATION OUT parameter list. An element is as defined by the SCSI-3 Medium Changer Commands (SMC) standard. The Element scope is optional for all device servers that implement PERSISTENT RESERVATION OUT.	3h – Fh	Reserved	Reserved
Code	Name	Description																
0h	LU	Logical Unit. Persistent reservation applies to the full logical unit. The LU scope is therefore implemented by all device servers that implement PERSISTENT RESERVATION OUT.																
1h	Extent	Persistent reservation applies to the specified extent. When Extent is the scope, it indicates that the persistent reservation applies to the extent of the logical unit defined by the Scope-Specific Address and Extent Length fields in the PERSISTENT RESERVATION OUT command's parameter list. Note that an extent is defined only for devices defining contiguous logical block addresses. The Extent scope is optional for all device servers that implement PERSISTENT RESERVATION OUT.																
2h	Element	Persistent reservation applies to the specified element. When Element is the scope, it indicates that the persistent reservation applies to the element of the logical unit defined by the Scope-Specific Address field in the PERSISTENT RESERVATION OUT parameter list. An element is as defined by the SCSI-3 Medium Changer Commands (SMC) standard. The Element scope is optional for all device servers that implement PERSISTENT RESERVATION OUT.																
3h – Fh	Reserved	Reserved																
Type	<p>The value of the Type field specifies the characteristics of the persistent reservation being established for all data blocks within the extent or within the logical unit. Refer to Table 5-56 for the applicable Type codes and their meanings</p>																	

PERSISTENT RESERVATION IN Command (5Eh) (continued)

Table 5-56 presents the definitions of the characters of the available “Type” values from the Type field of the PERSISTENT RESERVATION IN Read Reservations parameters.

Each of the codes provides handling instructions for READ operations, for WRITE operations, and for subsequent attempts to establish persistent reservations (referred to as “Additional Reservations Allowed” in the table).

Table 5-56 Persistent Reservation Type Codes and Their Meanings

Code	Name	Description
0h	READ Shared	<p>READS: Shared; any application client on any initiator may execute commands that perform transfers from the disk to the initiator.</p> <p>WRITES: Prohibited; any command from any initiator that performs a transfer from the initiator to the disk results in a reservation conflict.</p> <p>ADDITIONAL RESERVATIONS: Allowed; any initiator may reserve the logical unit, extents, or elements as long as the persistent reservations do not conflict with any reservations already known to the device server.</p>
1h	WRITE Exclusive	<p>READS: Shared; any application client on any initiator may execute commands that perform transfers from the disk to the initiator.</p> <p>WRITES: Exclusive; any command from any initiator other than the initiator that holds the persistent reservation that attempts a transfer to the disk results in a reservation conflict</p> <p>ADDITIONAL RESERVATIONS: Allowed; any initiator may reserve the logical unit, extents, or elements as long as the persistent reservations do not conflict with any reservations already known to the device server.</p>
2h	READ Exclusive	<p>READS: Exclusive; any command from any initiator other than the initiator that holds the persistent reservation that attempts a transfer from the disk results in a reservation conflict</p> <p>WRITES: Shared; any application client on any initiator may execute commands that perform transfers to the disk from the initiator.</p> <p>ADDITIONAL RESERVATIONS: Allowed; any initiator may reserve the logical unit, extents, or elements as long as the persistent reservations do not conflict with any reservations already known to the device server.</p>

(continued)

PERSISTENT RESERVATION IN Command (5Eh) (continued)

Table 5-56 Persistent Reservation Type Codes and Their Meanings (continued)

Code	Name	Description
3h	Exclusive Access	<p>READS: Exclusive; any command from any initiator other than the initiator holding the persistent reservation that attempts a transfer from the disk results in a reservation conflict.</p> <p>WRITES: Exclusive; any command from any initiator other than the initiator holding the persistent reservation that attempts a transfer to the disk results in a reservation conflict.</p> <p>ADDITIONAL RESERVATIONS: Restricted; any PERSISTENT RESERVATION OUT command with the Reserve service action from any initiator other than the initiator holding the persistent reservation results in a reservation conflict. The initiator that holds the persistent reservation can reserve the logical unit, extents, or elements as long as the persistent reservations do not conflict with any reservations already known to the device server.</p>
4h	Shared Access	<p>READS: Shared; any application client on any initiator may execute commands that perform transfers from the disk to the initiator.</p> <p>WRITES: Shared; ; any application client on any initiator may execute commands that perform transfers to the disk from the initiator.</p> <p>ADDITIONAL RESERVATIONS: Restricted; any PERSISTENT RESERVATION OUT command with the Reserve service action from any initiator other than the initiator holding the persistent reservation results in a reservation conflict. The initiator that holds the persistent reservation can reserve the logical unit, extents, or elements as long as the persistent reservations do not conflict with any reservations already known to the device server.</p>
5h	WRITE Exclusive, Registrants Only	<p>READS: Shared; any application client on any initiator may execute commands that perform transfers from the disk to the initiator.</p> <p>WRITES: Exclusive; any command from an initiator that has not previously performed a Register service action with the device server that attempts a transfer to the disk results in a reservation conflict</p> <p>ADDITIONAL RESERVATIONS: Allowed; any initiator may reserve the logical unit, extents, or elements as long as the persistent reservations do not conflict with any reservations already known to the device server.</p>

(continued)

*PERSISTENT RESERVATION IN Command (5Eh) (continued)**Table 5-56 Persistent Reservation Type Codes and Their Meanings (continued)*

Code	Name	Description
6h	Exclusive	<p>READS: Exclusive; any command from an initiator that has not previously performed a Register service action with the device server that attempts a transfer from the disk results in a reservation conflict.</p> <p>WRITES: Exclusive; any command from an initiator that has not previously performed a Register service action with the device server that attempts a transfer to the disk results in a reservation conflict</p> <p>ADDITIONAL RESERVATIONS: Allowed; any initiator may reserve the logical unit, extents, or elements as long as the persistent reservations to not conflict with any reservations already known to the device server.</p>
	Registrants Only	
7h - Fh	Reserved	N/A

PERSISTENT RESERVATION IN Command (5Eh) (continued)

Table 5-57 represents graphically all possible combinations of Persistent Reservations Being Attempted when Persistent Reservations are already held for each of the types of persistent reservations.

Table 5-57 When Do Conflicts Between Existing Reservations and New Reservations Exist?

Persistent Reservation Already Held														
Persistent Reservation Being Attempted	READ Shared		WRITE Exclusive		READ Exclusive		Exclusive Access *		Shared Access *		WRITE Exclusive RO		Exclusive Access RO	
	LU	EX	LU	EX	LU	EX	LU	EX	LU	EX	LU	EX	LU	EX
	READ Shared LU	N	N	Y	Y	Y	Y	Y	Y	N	N	Y	Y	Y
READ Shared EX	N	N	Y	O	Y	O	Y	O	N	N	Y	O	Y	O
WRITE Exclusive LU	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
WRITE Exclusive EX	Y	O	Y	O	Y	O	Y	O	Y	O	Y	O	Y	O
READ Exclusive LU	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
READ Exclusive EX	Y	O	Y	O	Y	O	Y	O	Y	O	Y	O	Y	O
Exclusive Access * LU	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Exclusive Access * EX	Y	O	Y	O	Y	O	Y	O	Y	O	Y	O	Y	O
Shared Access * LU	N	N	Y	Y	Y	Y	Y	Y	N	N	Y	Y	Y	Y
Shared Access * EX	N	N	Y	O	Y	O	Y	O	N	N	Y	O	Y	O
WRITE Exclusive RO LU	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	N	N	N	N
WRITE Exclusive RO EX	Y	O	Y	O	Y	O	Y	O	O	O	N	N	N	N
Exclusive Access RO LU	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	N	N	N	N
Exclusive Access RO EX	Y	O	Y	O	Y	O	Y	O	Y	O	N	N	N	N

KEY: LU = Logical Unit Scope
 EX = Extent or Element Scope
 RO = Registrants Only
 * = Conflicts with all reservation requests from other initiators.

N = No Conflict
 Y = Conflict
 O = Conflict occurs if extent or element overlaps with existing extent or element reservation.

5.19 PERSISTENT RESERVATION OUT Command (5Fh)

The PERSISTENT RESERVATION OUT command is a 10-byte command used to reserve a logical unit or an extent within a logical unit for the exclusive or shared use by an initiator. The command is used in conjunction with the PERSISTENT RESERVATION IN command; it is not used with the RESERVE and RELEASE commands.

Persistent reservations conflict with reservations made via the RESERVE command. Initiators that perform PERSISTENT RESERVATION OUT actions are identified by a reservation key assigned by the application client. The client may use the PERSISTENT RESERVATION IN command to identify which other initiators within a system hold conflicting or invalid persistent reservations and use the PERSISTENT RESERVATION OUT command to preempt those reservations if necessary.

Note that since persistent reservations are not reset by the TARGET RESET task management function or other global actions, they can be used to enact device sharing among multiple initiators. The PERSISTENT RESERVATION OUT and PERSISTENT RESERVATION IN commands provide the means for resolving contentions in multiple-initiator systems with multiple port target. By using the reservation key to identify persistent reservations, it is possible to determine which ports hold conflicting persistent reservations and to take over such reservations from failing or “greedy” initiators.

The figure below illustrates the format of the PERSISTENT RESERVATION OUT command; the table that follows explains the data fields of the command.

Bit	7	6	5	4	3	2	1	0
Byte								
0	Operation Code (5Fh)							
1	Reserved			Service Action				
2	Scope				Type			
3 - 6	Reserved							
7 - 8	(MSB) Parameter List Length (18h)							(LSB)
9	Control							

Figure 5-54 PERSISTENT RESERVATION OUT Command Descriptor Block — Data Format

*PERSISTENT RESERVATION OUT Command (5Fh) (continued)**Table 5-58 PERSISTENT RESERVATION OUT Command — Field Descriptions*

Data Field	Description	
Service Action	Service actions that require information about persistent reservation and registrations may require enabling of nonvolatile memory within the logical unit. Service action codes available are:	
	Code	Name Description
	00h	Register Register a reservation key with the device server
	01h	Reserve Create a persistent reservation using a reservation key
	02h	Release Release a persistent reservation
	03h	Clear Clear all reservation keys and all persistent reservations
	04h	Pre-empt Pre-empt persistent reservations from another initiator
	05h	Pre-empt & Clear Pre-empt persistent reservations from another initiator and clear the task set for the pre-empted initiator
	06 – 1Fh	Reserved Reserved
	Refer to Table 5-59 for detailed descriptions of each of the service action codes.	

(continued)

*PERSISTENT RESERVATION OUT Command (5Fh) (continued)**Table 5-58 PERSISTENT RESERVATION OUT Command — Field Descriptions (continued)*

Data Field	Description		
Scope	The value in this field indicates whether a persistent reservation applies to an entire logical unit, to a part of the logical unit (defined as an extent), or to an element. The values for the Scope field are:		
	Code	Name	Description
	0h	LU	Logical Unit. Persistent reservation applies to the full logical unit. The LU scope is therefore implemented by all device servers that implement PERSISTENT RESERVATION OUT.
	1h	Extent	Persistent reservation applies to the specified extent. When Extent is the scope, it indicates that the persistent reservation applies to the extent of the logical unit defined by the Scope-Specific Address and Extent Length fields in the PERSISTENT RESERVATION OUT command's parameter list. Note that an extent is defined only for devices defining contiguous logical block addresses. The Extent scope is optional for all device servers that implement PERSISTENT RESERVATION OUT.
	2h	Element	Persistent reservation applies to the specified element. When Element is the scope, it indicates that the persistent reservation applies to the element of the logical unit defined by the Scope-Specific Address field in the PERSISTENT RESERVATION OUT parameter list. An element is as defined by the SCSI-3 Medium Changer Commands (SMC) standard. The Element scope is optional for all device servers that implement PERSISTENT RESERVATION OUT.
	3h – Fh	Reserved	Reserved

(continued)

*PERSISTENT RESERVATION OUT Command (5Fh) (continued)**Table 5-58 PERSISTENT RESERVATION OUT Command — Field Descriptions (continued)*

Data Field	Description
Type	The value of the Type field specifies the characteristics of the persistent reservation being established for all data blocks within the extent or within the logical unit. Refer to Table 5-60 for the applicable Type codes and their meanings
Parameter List Length	Fields contained in the PERSISTENT RESERVATION OUT parameter list specify the reservation keys and extent information required to perform a persistent reservation service action. The parameter list is 24 bytes in length; the Parameter List Length field contains 24 (18h) bytes.

Table 5-59 provides detailed descriptions of each of the PERSISTENT RESERVATION OUT command's seven possible service actions (Service Action codes appear in bits 0 – 4 of Byte 1).

Table 5-59 PERSISTENT RESERVATION OUT Command's Service Action Descriptions

Code	Name	Description
00h	Register	<p>When the command executes a Register service action, it registers a reservation key with a device server without generating a reservation. The device server holds these reservation keys from each initiator that performs a PERSISTENT RESERVATION OUT command with a Register service action until the key is changed by a new PERSISTENT RESERVATION OUT command with Register service action from the same initiator, or until the initiator registration is removed by:</p> <ul style="list-style-type: none"> • Powering down the logical unit, if the last Activate Persist Through Power Loss (APTPL; see Figure 5-55 and Table 5-61) received by the device server was 0; • Performing a Clear service action; • Performing a Pre-empt service action; • Performing a Pre-empt and Clear service action; or • Performing a Register service action from the same initiator with the value of the service action reservation key set to 0. <p>When a reservation key has not yet been established or when the reservation key has been removed, a reservation key of 0 is used when the initiator performs a PERSISTENT RESERVATION OUT with the Register service action. When the reservation has been removed, no information is reported for the initiator in the Read Keys service action of the resulting PERSISTENT RESERVATION IN command.</p>

(continued)

*PERSISTENT RESERVATION OUT Command (5Fh) (continued)**Table 5-59 PERSISTENT RESERVATION OUT Command's Service Action Descriptions (continued)*

Code	Name	Description
01h	Reserve	<p>A PERSISTENT RESERVATION OUT command with Reserve service action creates a persistent reservation with a specified Scope and Type.</p> <p>Persistent reservations are not superseded by a new persistent reservation from any initiator except by the execution of a PERSISTENT RESERVATION OUT command that specifies either a Pre-empt or Pre-empt and Clear service action. New persistent reservations that do not conflict with an existing persistent reservation execute normally. Persistent reservations of logical units or extents having the same Type value are permitted if no conflicting persistent reservations are held by another initiator. When these types of overlapping reservations are released, each of the extent reservations and logical unit reservations are removed with a separate Release service action. Multiple identical reservations from the same initiator are all released simultaneously via a single Release service action that matches the reservations.</p>
02h	Release	<p>A PERSISTENT RESERVATION OUT command with Release service action removes a persistent reservation held by the same initiator.</p> <p>The fields associated with a Release service action match fields of the active persistent reservation. Sending of a PERSISTENT RESERVATION OUT command that specifies a Release service action when no persistent reservation exists from that initiator does not result in an error. Instead, the device server returns a GOOD message without altering any other reservation: the reservation key is not changed by the Release service action.</p> <p>The device server returns a CHECK CONDITION status for any PERSISTENT RESERVATION OUT command that specifies the release of a persistent reservation held by the requesting initiator that matches some but not all of the Scope, Type, Reservation Key, and extent values. The sense key is set to ILLEGAL REQUEST and additional sense data is set to INVALID RELEASE OF ACTIVE PERSISTENT RESERVATION. Attempts to release persistent reservations in which none of the Scope, Type, Reservation Key, and extent values match an existing persistent reservation held by the initiator making the request are not errors.</p> <p>An active persistent reservation may also be released by:</p> <ul style="list-style-type: none"> Powering off. When the most recent APTPL value received by the device server is 0, a power-off performs a hard reset, clears all persistent reservations, and removes all registered reservation keys; <p>or</p> <ul style="list-style-type: none"> Executing a PERSISTENT RESERVATION OUT command from another initiator with a persistent reserve service action of Pre-empt or Pre-empt and Clear. <p>Note that a Release service action should not be performed if any operations interlocked by the persistent reservation have not yet completed.</p>

(continued)

*PERSISTENT RESERVATION OUT Command (5Fh) (continued)**Table 5-59 PERSISTENT RESERVATION OUT Command's Service Action Descriptions (continued)*

Code	Name	Description
03h	Clear	<p>A PERSISTENT RESERVATION OUT command with a successful Clear service action removes all persistent reservations for all initiators. All reservation keys are also removed. Any commands from any initiator that have been accepted by the device server as non-conflicting continue their normal executions.</p> <p>A UNIT ATTENTION condition is established for all initiators for the cleared logical unit. The sense key is set to UNIT ATTENTION; the additional sense data is set to RESERVATIONS PREEMPTED.</p> <p>Note that applications should not use the Clear action service except during recoveries associated with initiator or system reconfiguration, since data integrity may be compromised.</p>
04h	Pre-empt	<p>A PERSISTENT RESERVATION OUT command with a successful Pre-empt service action removes all persistent reservations for all initiators that have been registered with the Service action Reservation key specified in the PERSISTENT RESERVATION OUT command's parameter list. A persistent reservation is also established for the pre-empting initiator. Any commands from any initiator that have been accepted by the device server as non-conflicting continue their normal executions. If a PERSISTENT RESERVATION OUT command is sent that specifies a Pre-empt service action and no persistent reservation exists for the initiator identified by the Service action Reservation key, it is not an error condition.</p> <p>A UNIT ATTENTION condition is established for the pre-empted initiators. The sense key is set to UNIT ATTENTION; the additional sense data is set to RESERVATIONS PREEMPTED. Commands that follow are subject to the persistent reservation restrictions set by the pre-empting initiator.</p> <p>The persistent reservation thus created by the pre-empting initiator is defined by the Scope and Type fields of the PERSISTENT RESERVATION OUT command and the corresponding fields of the command's parameter list.</p> <p>The registration keys for the pre-empted initiators are removed by the Pre-empt service action; the reservation key for an initiator that has performed a Pre-empt service action with its own Reservation key specified in the Service action Reservation key remains unchanged, although all other specified releasing actions and reservation actions are performed.</p> <p>Note that persistent reservations are not superseded by a new persistent reservation from any initiator except by the execution of a PERSISTENT RESERVATION OUT that specifies either the Pre-empt or the Pre-empt and Clear service actions. New persistent reservations that do not conflict with an existing persistent reservation execute normally. The persistent reservation of a logical unit or extents having the same Type value are permitted if no conflicting persistent reservations other than the reservations being pre-empted are held by another initiator.</p>

(continued)

*PERSISTENT RESERVATION OUT Command (5Fh) (continued)**Table 5-59 PERSISTENT RESERVATION OUT Command's Service Action Descriptions (continued)*

Code	Name	Description
05h	Pre-empt & Clear	<p>A PERSISTENT RESERVATION OUT command with a Pre-empt & Clear service action removes all persistent reservations for all initiators that have been registered with the Service action Reservation key specified in the PERSISTENT RESERVATION OUT command's parameter list. It also establishes a persistent reservation for the pre-empting initiator. Any commands from the initiators being pre-empted are terminated as if an ABORT TASK management function had been performed by the pre-empted initiator. If a PERSISTENT RESERVATION OUT command is sent that specifies a Pre-empt & Clear service action and no persistent reservation exists for the initiator identified by the Service action Reservation key, it is not an error condition. If the key is registered, however, the Clear portion of the action executes normally.</p> <p>A UNIT ATTENTION condition is established for the pre-empted initiators. The sense key is set to UNIT ATTENTION; the additional sense data is set to RESERVATIONS PREEMPTED. Commands that follow, and retries of commands that timed out because there were cleared are subject to the persistent reservation restrictions set by the pre-empting initiator.</p> <p>The persistent reservation thus created by the pre-empting initiator is defined by the Scope and Type fields of the PERSISTENT RESERVATION OUT command and the corresponding fields of the command's parameter list.</p> <p>The Pre-empt & Clear service action clears any ACA or CA condition associated with the initiator that is pre-empted and clears any tasks with an ACA attribute from that initiator.</p> <p>Any Asynchronous Event Reporting operations in progress that were initiated by the device server are unaffected by a Pre-empt and Clear service action.</p> <p>The reservation key for the other initiators pre-empted are removed by the Pre-empt & Clear service action. The reservation key for an initiator that has sent a Pre-empt & Clear action with its own reservation key specified in the service action's reservation key remains unchanged, although all other specified clearing actions, releasing actions, and reservation actions are performed.</p> <p>Persistent reservations are not superseded by a new persistent reservation from any initiator except via execution of a PERSISTENT RESERVATION OUT that specifies either the Pre-empt or Pre-empt & Clear service action. New persistent reservations not in conflict with an existing persistent reservation execute normally. The persistent reservation of a logical unit or extent that has the same Type value is permitted as long as no conflicting persistent reservations other than the reservations being pre-empted are permitted.</p>

PERSISTENT RESERVATION OUT Command (5Fh) (continued)

Table 5-60 presents the definitions of the characters of the available “Type” values from the Type field of the PERSISTENT RESERVATION IN Read Reservations parameters.

Each of the codes provides handling instructions for READ operations, for WRITE operations, and for subsequent attempts to establish persistent reservations (referred to as “Additional Reservations Allowed” in the table).

Table 5-60 Persistent Reservation Type Codes and Their Meanings

Code	Name	Description
0h	READ Shared	<p>READS: Shared; any application client on any initiator may execute commands that perform transfers from the disk to the initiator.</p> <p>WRITES: Prohibited; any command from any initiator that performs a transfer from the initiator to the disk results in a reservation conflict.</p> <p>ADDITIONAL RESERVATIONS: Allowed; any initiator may reserve the logical unit, extents, or elements as long as the persistent reservations do not conflict with any reservations already known to the device server.</p>
1h	WRITE Exclusive	<p>READS: Shared; any application client on any initiator may execute commands that perform transfers from the disk to the initiator.</p> <p>WRITES: Exclusive; any command from any initiator other than the initiator that holds the persistent reservation that attempts a transfer to the disk results in a reservation conflict</p> <p>ADDITIONAL RESERVATIONS: Allowed; any initiator may reserve the logical unit, extents, or elements as long as the persistent reservations do not conflict with any reservations already known to the device server.</p>
2h	READ Exclusive	<p>READS: Exclusive; any command from any initiator other than the initiator that holds the persistent reservation that attempts a transfer from the disk results in a reservation conflict</p> <p>WRITES: Shared; any application client on any initiator may execute commands that perform transfers to the disk from the initiator.</p> <p>ADDITIONAL RESERVATIONS: Allowed; any initiator may reserve the logical unit, extents, or elements as long as the persistent reservations do not conflict with any reservations already known to the device server.</p>

(continued)

PERSISTENT RESERVATION OUT Command (5Eh) (continued)
Table 5-60 Persistent Reservation Type Codes and Their Meanings (continued)

Code	Name	Description
3h	Exclusive Access	<p>READS: Exclusive; any command from any initiator other than the initiator holding the persistent reservation that attempts a transfer from the disk results in a reservation conflict.</p> <p>WRITES: Exclusive; any command from any initiator other than the initiator holding the persistent reservation that attempts a transfer to the disk results in a reservation conflict.</p> <p>ADDITIONAL RESERVATIONS: Restricted; any PERSISTENT RESERVATION OUT command with the Reserve service action from any initiator other than the initiator holding the persistent reservation results in a reservation conflict. The initiator that holds the persistent reservation can reserve the logical unit, extents, or elements as long as the persistent reservations do not conflict with any reservations already known to the device server.</p>
4h	Shared Access	<p>READS: Shared; any application client on any initiator may execute commands that perform transfers from the disk to the initiator.</p> <p>WRITES: Shared; ; any application client on any initiator may execute commands that perform transfers to the disk from the initiator.</p> <p>ADDITIONAL RESERVATIONS: Restricted; any PERSISTENT RESERVATION OUT command with the Reserve service action from any initiator other than the initiator holding the persistent reservation results in a reservation conflict. The initiator that holds the persistent reservation can reserve the logical unit, extents, or elements as long as the persistent reservations do not conflict with any reservations already known to the device server.</p>
5h	WRITE Exclusive, Registrants Only	<p>READS: Shared; any application client on any initiator may execute commands that perform transfers from the disk to the initiator.</p> <p>WRITES: Exclusive; any command from an initiator that has not previously performed a Register service action with the device server that attempts a transfer to the disk results in a reservation conflict</p> <p>ADDITIONAL RESERVATIONS: Allowed; any initiator may reserve the logical unit, extents, or elements as long as the persistent reservations do not conflict with any reservations already known to the device server.</p>

(continued)

PERSISTENT RESERVATION OUT Command (5Eh) (continued)

Table 5-60 Persistent Reservation Type Codes and Their Meanings (continued)

Code	Name	Description
6h	Exclusive Registrants Only	<p>READS: Exclusive; any command from an initiator that has not previously performed a Register service action with the device server that attempts a transfer from the disk results in a reservation conflict.</p> <p>WRITES: Exclusive; any command from an initiator that has not previously performed a Register service action with the device server that attempts a transfer to the disk results in a reservation conflict</p> <p>ADDITIONAL RESERVATIONS: Allowed; any initiator may reserve the logical unit, extents, or elements as long as the persistent reservations to not conflict with any reservations already known to the device server.</p>
7h - Fh	Reserved	N/A

The PERSISTENT RESERVATION OUT command requires a parameter list, illustrated in the following figure and defined in the following table. Each of the fields of the parameter list are sent for every PERSISTENT RESERVATION OUT command, even if the field is not required for the specific Service action and/or Scope values.

Bit	7	6	5	4	3	2	1	0
Byte								
0 - 7	(MSB) Reservation Key							(LSB)
8 - 15	(MSB) Service Action Reservation Key							(LSB)
16 - 19	(MSB) Scope-Specific Address							(LSB)
20	Reserved						APTPL	
21	Reserved							
22 - 23	(MSB) Extent Length							(LSB)

Figure 5-55 PERSISTENT RESERVATION OUT Parameter List — Data Format

PERSISTENT RESERVATION OUT Command (5Eh) (continued)
Table 5-61 PERSISTENT RESERVATION OUT Parameter List — Field Descriptions

Data Field	Description
Reservation Key	<p>This field contains an 8-byte token that is provided by the application client to the device server to identify which initiator is the source of the PERSISTENT RESERVATION OUT command. The device server verifies that the Reservation Key in the PERSISTENT RESERVATION OUT command matches the Reservation Key that is registered for the initiator from which the command is received. If there is no match, the device server returns a RESERVATION CONFLICT status. The Reservation Key of the initiator is valid for all Service action and Scope values.</p>
Service Action Reservation Key	<p>This field contains information needed for 3 service actions: the Register service action, the Pre-empt service action, and the Pre-empt & Clear service action. The Service Action Reservation Key is ignored for all other service actions.</p> <p>For the Register service action, the Service Action Reservation Key field contains the new Reservation Key to be registered.</p> <p>For the Pre-empt and the Pre-empt & Clear service actions, the Service Action Reservation Key contains the reservation key of the persistent reservations that are being pre-empted. For the Pre-empt and the Pre-empt & Clear actions, any failure of the Service Action Reservation Key to match any registered keys results in the device server returning a RESERVATION CONFLICT status.</p>
Scope-Specific Address	<p>If the Scope is an Extent reservation, this field contains the Logical Block Address of the extent and the Extent Length field (bytes 22 – 23) contain the number of blocks in the extent.</p> <p>If the Scope signifies an Element reservation, the Scope-Specific Address field contains the Element Address, zero-filled in the most significant bytes to fit the field; the Extent Length is set to 0.</p> <p>If the Service action is Register or Clear, or if the Scope is a Logical Unit reservation, both the Scope-Specific Address and Extent Length fields are set to 0.</p>
APTPL	<p>Activate Persist Through Power Loss. This bit is valid only for Register service actions; it is ignored for all other types of service actions.</p> <p>Support for APTPL set to 1 is optional. If a device server that does not support APTPL receives a 1 in that bit in a Register service action, it returns a CHECK CONDITION status, with the sense key set to ILLEGAL REQUEST and additional sense data set to INVALID FIELD IN PARAMETER LIST.</p> <p>If the last valid APTPL bit value received by the device server is 0, the loss of power in the target releases any persistent reservations and removes all reservation keys. If the last valid APTPL bit value is 1, the logical unit retains all persistent reservations and all reservation keys for all initiators even if power is lost and later returned. The most recently received valid APTPL value from any initiator governs the logical unit's behavior in the event of a power loss.</p>

PERSISTENT RESERVATION OUT Command (5Eh) (continued)

The table below illustrates which fields are set by the application client and interpreted by the device server for each Service and Scope value.

Table 5-62 Device Server Interpretation of Service and Scope Value

Service Action	Allowed Scope	Parameters		
		Type	Service Action Reservation Key	Element or Element Parameters
Register	Ignored	Ignored	Valid	Ignored
Reserve	LU	Valid	Ignored	Ignored
Reserve	Extent	Valid	Ignored	Extent Valid
Reserved	Element	Valid	Ignored	Element Valid
Release	LU	Valid	Ignored	Ignored
Release	Extent	Valid	Ignored	Extent Valid
Release	Element	Valid	Ignored	Element Valid
Clear	Ignored	Ignored	Ignored	Ignored
Pre-empt	LU	Valid	Valid	Ignored
Pre-empt	Extent	Valid	Valid	Extent Valid
Pre-empt	Element	Valid	Valid	Element Valid
Pre-empt & Clear	LU	Valid	Valid	Ignored
Pre-empt & Clear	Extent	Valid	Valid	Extent Valid
Pre-empt & Clear	Element	Valid	Valid	Element Valid

5.20 READ (6) Command (08h)

The READ (6) command requests that data be transferred from the drive to the initiator. The amount of data transferred is specified by the transfer length field. The starting point is specified by the logical sector address.

Bit	7	6	5	4	3	2	1	0
Byte								
0	Operation Code (08h)							
1	Reserved			Logical Block Address				
2 - 3	Logical Block Address (continued)							
4	Transfer Length							
5	Control							

Figure 5-56 READ (6) Command – Data Format

Table 5-63 READ (6) Command — Field Descriptions

Field	Description
Logical Block Address	Specifies the logical block address where the read operation will begin.
Transfer Length	Specifies the number of contiguous logical blocks of data to transfer. A transfer length of 0 indicates that 256 logical blocks will be transferred. Any other value indicates the number of logical blocks that will be transferred.

5.21 READ (10) Command (28h)

The READ (10) command requests that data be transferred from the drive to the initiator. The amount of data transferred is specified by the transfer length field. The starting point is specified by the logical sector address.

Bit	7	6	5	4	3	2	1	0
Byte								
0	Operation Code (28h)							
1	Reserved			DPO	FUA	Reserved		RelAdr
2 - 5	Logical Block Address							
6	Reserved							
7 - 8	Transfer Length							
9	Control							

Figure 5-57 READ (10) Command — Data Format

Table 5-64 READ (10) Command — Field Descriptions

Field	Description
DPO	Disable Page Out. When the value is 0, the drive caches the data read, using its normal Logical Read Unit policies. When the value is 1, the drive assigns to logical blocks accessed by the READ (10) command the lowest priority for being fetched into or retained by the cache.
FUA	Force Unit Access. When the value is 0, the drive is allowed to satisfy the read with cache-resident data. A value of 1 indicates that the drive accesses the media in performing the command before returning GOOD status.
RelAdr	Relative Addressing. The drive does not support relative addressing. Must be 0.
Logical Block Address	Specifies the logical sector address where the read operation will begin.
Transfer Length	Specifies the number of contiguous logical blocks of data to transfer. A transfer length of 0 indicates that no logical blocks are to be transferred. Any other value indicates the number of logical blocks to be transferred. The drive returns a CHECK CONDITION status with the correct sense data if the specified transfer extends beyond the logical sector space.

5.22 READ BUFFER Command (3Ch)

The READ BUFFER command is used in conjunction with the WRITE BUFFER command. It allows the initiator to gain access to the data storage areas within the drive's controller. In particular, these two commands are used to check the integrity of the buffer-initiator data path, independent of the media, typically for fault isolation in the event of a failure. In addition, these commands can be used to verify the caching policies of the drive controller or other kinds of diagnostic and verification work.

Both the READ BUFFER and the WRITE BUFFER commands have a buffer ID and a buffer offset data field defined in their Control Descriptor Block. In the Combined Header and Data Mode, if both are set to 0, reference is made to a single 512-byte buffer set aside in the drive controller's data memory (independent of cache). This allows an initiator to verify the data path from the controller's data memory to the initiator without affecting the drive's internal data caches.

Bit	7	6	5	4	3	2	1	0
Byte								
0	Operation Code (3Ch)							
1	Reserved				Mode			
2	Buffer ID							
3 - 5	Buffer Offset							
6 - 8	Allocation Length							
9	Control							

Figure 5-58 READ BUFFER Command — Data Format

READ BUFFER Command (3Ch) (continued)

Table 5-65 READ BUFFER Command — Field Descriptions

Field	Description
Mode	Four modes are supported by the drive; Combined Header and Data, Data Only, Descriptor Only, and Read Data from Echo Buffer. The permissible values for Buffer ID, Buffer Offset, are as follows:
0000b	<p>Combined Header and Data. – In this mode, data to be transferred is preceded by a 4-byte header that contains the buffer capacity. The resulting descriptor and data are transferred to the initiator during the command's data-in buffer transfer.</p> <p>Buffer ID must be 0</p> <p>Buffer Offset must be 0</p> <p>Allocation length must be set to between 0 and 516 (4 bytes of header and up to 512 bytes of data).</p>
0010b	<p>Data Only.– In this mode, the data-in buffer transfer contains data from the controller's memory.</p> <p>Buffer ID. When non-zero, the value specifies a specific track line in the controller's cache. If Buffer ID of zero is specified, the buffer referenced is the same as if a Buffer ID of one were specified.</p> <p>Buffer Offset. When non-zero, the value of Buffer Offset can range from 0 to one less than the number of sectors/track multiplied by 512. This value must be either a zero or a multiple of 512.</p> <p>Allocation Length Must be between zero and the maximum number of sectors/track multiplied by 512.</p>
0011b	<p>Descriptor Only</p> <p>When requesting a descriptor (a description of the buffer referenced in the Buffer ID field), the resulting data is the offset boundary requirements and the length of each buffer.</p> <p>Buffer ID When non-zero, the value specifies a specific track line within the controller's cache. If Buffer ID of 0 is specified, the buffer referenced is the same as if a Buffer ID of 1 were specified.</p> <p>Buffer Offset Must be zero</p> <p>Allocation Length Must be between 4 (header) and the maximum number of sectors/track multiplied by 512 + header size.</p>
1010b	<p>Read Data from Echo Buffer</p> <p>In this mode, the drive transfers data to the host from the echo buffer. The echo buffer will transfer the same data as when the WRITE BUFFER command with the mode field set to echo buffer was issued.</p> <p>Buffer ID Ignored.</p> <p>Buffer Offset Ignored.</p>

5.23 READ CAPACITY Command (25h)

The READ CAPACITY command reports formatted capacity and track-boundary information.

Bit	7	6	5	4	3	2	1	0
Byte								
0	Operation Code (25h)							
1	Reserved							RelAdr
2 - 5	Logical Block Address							
6 - 7	Reserved							
8	Reserved							PMI
9	Control							

Figure 5-59 READ CAPACITY Command — Data Format

Table 5-66 READ CAPACITY Command — Field Descriptions

Field	Description
RelAdr	Relative Addressing is not supported by the drive; the field must be 0.
Logical Block Address	Used in conjunction with the PMI bit. Must be 0 when the PMI bit is 0.
PMI	<p>Partial Media Indicator. When the value in the field is 0, the Returned Logical Block Address (Figure 5-60) is that of the last logical sector of the logical unit.</p> <p>When the value is 1, the Returned Logical Block Address is that of the last sector on the same cylinder as the Logical Block Address supplied in the Command Descriptor Block's LBA field. That is, the Returned Logical Block Address is that of the Logical Block Address before a substantial delay in data transfer will be encountered).</p> <p>In either case, Block Length in Bytes (Figure 5-60) is the unit's logical sector size, as determined by the Block Length field in the Mode parameter Block Descriptor.</p>

READ CAPACITY Command (25 h) (continued)

Figure 5-60 shows the READ CAPACITY Returned Data - Data Format that shall be sent during the data-in buffer transfer for the command. Refer to the PMI description in Table 5-66 for a description of these data fields.

Bit Byte	7	6	5	4	3	2	1	0
0 - 3	Returned Logical Block Address							
4 - 7	Block Length in Bytes							

Figure 5-60 READ CAPACITY Returned Data — Data Format

5.24 READ DEFECT DATA Command (10) (37h)

The READ DEFECT DATA (10) command directs the drive to report defect information about the drive's medium. Refer to the FORMAT UNIT command for information about defects. The data sent to the host in the command's data-in buffer transfer consists of a four-byte Defect List header (Figure 5-64), followed by any Defect Descriptors (Figures 5-62 and 5-64).

The READ DEFECT DATA (10) Command Descriptor Block is shown in Figure 5-66; the data fields are described in Table 5-67.

Bit	7	6	5	4	3	2	1	0
Byte								
0	Operation Code (37h)							
1	Reserved							
2	Reserved			Plist	Glist	Defect List Format		
3 - 6	Reserved							
7 - 8	(MSB) Allocation Length (LSB)							
9	Control							

Figure 5-61 READ DEFECT DATA (10) Command — Data Format

Table 5-67 READ DEFECT DATA (10) Command — Field Description

Field	Description
PList	When the value is 0, the returned data does not include the Primary Defect List. When the value is 1, the returned data includes the Primary Defect List. The defects are returned in ascending order.
GList	When the value is 0, the returned data does not include the Grown Defect List. When the value is 1, the returned data includes the Grown Defect List. The defects are returned in order of occurrence. When both PList and GList are set to 1, the Primary Defect List and the Grown Defect List are returned. The PList is returned followed by the GList (the lists are not merged). When both are set to 0, only the Defect List Header (Figure 5-64) is reported to the initiator.

(continued)

READ DEFECT DATA Command (10) (37h) (continued)

Table 5-67 READ DEFECT DATA (10) Command — Field Descriptions (continued)

Field	Description
Defect List Format	This field specifies the format in which the defect descriptors are to be returned. The supported formats are:
000b	Block Format (Figure 5-62). ¹
100b	Bytes from the Index (Figure 5-63). The data returned from this mode does not include defects outside normal user space.
101b	Physical Sector Format (Figure 5-63). Default format. The defects returned from this mode do not include defects outside normal user space.
110b	Vendor-Unique Format. This format is valid only for the READ DEFECT DATA commands. The data returned from this mode does not include defects outside normal user space. Note that the vendor-unique format is proprietary in nature and thus is not documented in this Product Manual.

¹ All in-line spared blocks do not have an associated LBA, and, therefore, will not be seen in Logical Block Format.

Figure 5-62 shows the Block Format returned when the Defect List Format is 000b.

Bit	7	6	5	4	3	2	1	0
Byte								
0 - 3	Defective Block Address							

Figure 5-62 Defect Descriptor — Block Format

Figure 5-63 shows the Defective Bytes From Index Format when the Defect List Format is 100b or the Physical Sector Format when the Defect List Format is 101b.

Bit	7	6	5	4	3	2	1	0
Byte								
0 - 2	Cylinder Number of Defect							
3	Head Number of Defect							
4 - 7	Defective Bytes From Index / Defective Sector Number							

Figure 5-63 Defect Descriptor — Bytes From Index Format, or Physical Sector Format

READ DEFECT DATA Command (10) (37h) (continued)

Bit	7	6	5	4	3	2	1	0
Byte								
0	Reserved							
1	Reserved			Plist	Glist	Defect List Format		
2 - 3	Defect List Length							

*Figure 5-64 Defect List Header — Data Format**Table 5-68 Defect List Header — Field Descriptions*

Field	Description
PList	Primary Defect List. When the value is 0, the returned data does not include the Primary Defect List. When the value is 1, the returned data includes the Primary Defect List.
GList	Grown Defect List. When the value is 0, the returned data does not include the Grown Defect List. When the value is 1, the returned data includes the Grown Defect List.
Defect List Format	This field indicates the format of the descriptors returned. If this differs from the format requested in the Command Descriptor Block, a recovered error is returned.
Defect List Length	This indicates the actual number of bytes that follow the defect list header. The drive truncates the list if the allocation length is not sufficient to accommodate the entire list (CHECK CONDITION status is not returned when this happens). The initiator is responsible for comparing the defect list length and the allocation length to ensure that a partial list was not received.

5.25 READ DEFECT DATA Command (12) (B7h)

The READ DEFECT DATA (12) command directs the drive to report defect information about the drive's medium. Refer to the FORMAT UNIT command for information about defects. The data sent to the host in the command's data-in buffer transfer consists of a four-byte Defect List header (Figure 5-68), followed by any Defect Descriptors (Figures 5-66 and 5-67).

The READ DEFECT DATA (12) Command Descriptor Block is shown in Figure 5-65; the data fields are described in Table 5-69.

Bit	7	6	5	4	3	2	1	0
Byte								
0	Operation Code (B7h)							
1	Reserved			Plist	Glist	Defect List Format		
2 - 5	Reserved							
6 - 9	(MSB) Allocation Length (LSB)							
10	Reserved							
11	Control							

Figure 5-65 READ DEFECT DATA (12) Command — Data Format

Table 5-69 READ DEFECT DATA (12) Command — Field Description

Field	Description
PList	When the value is 0, the returned data does not include the Primary Defect List. When the value is 1, the returned data includes the Primary Defect List. The defects are returned in ascending order.
GList	When the value is 0, the returned data does not include the Grown Defect List. When the value is 1, the returned data includes the Grown Defect List. The defects are returned in order of occurrence. When both PList and GList are set to 1, the Primary Defect List and the Grown Defect List are returned. The Plist is returned followed by the Glist (the lists are not merged). When both are set to 0, only the Defect List Header (Figure 5-68) is reported to the initiator.

(continued)

READ DEFECT DATA Command (12) (B7h) (continued)

Table 5-69 READ DEFECT DATA (12) Command — Field Descriptions (continued)

Field	Description
Defect List Format	This field specifies the format in which the defect descriptors are to be returned. The supported formats are:
000b	Block Format (Figure 5-66). ¹
100b	Bytes from the Index (Figure 5-67). The data returned from this mode does not include defects outside normal user space.
101b	Physical Sector Format (Figure 5-67). Default format. The defects returned from this mode do not include defects outside normal user space.
110b	Vendor-Unique Format. This format is valid only for the READ DEFECT DATA commands. The data returned from this mode does not include defects outside normal user space. Note that the vendor-unique format is proprietary in nature and thus is not documented in this Product Manual.

¹ All in-line spared blocks do not have an associated LBA, and, therefore, will not be seen in Logical Block Format.

Figure 5-66 shows the Block Format returned when the Defect List Format is 000b.

Bit	7	6	5	4	3	2	1	0
Byte								
0 - 3	Defective Block Address							

Figure 5-66 Defect Descriptor — Block Format

Figure 5-67 shows the Defective Bytes From Index Format when the Defect List Format is 100b or the Physical Sector Format when the Defect List Format is 101b.

Bit	7	6	5	4	3	2	1	0
Byte								
0 - 2	Cylinder Number of Defect							
3	Head Number of Defect							
4 - 7	Defective Bytes From Index / Defective Sector Number							

Figure 5-67 Defect Descriptor — Bytes From Index Format, or Physical Sector Format

READ DEFECT DATA Command (12) (B7h) (continued)

Bit Byte	7	6	5	4	3	2	1	0
0	Reserved							
1	Reserved			PList	GList	Defect List Format		
2 - 3	Reserved							
4 - 7	Defect List Length							

Figure 5-68 Defect List Header — Data Format

Table 5-70 Defect List Header — Field Descriptions

Field	Description
PList	Primary Defect List. When the value is 0, the returned data does not include the Primary Defect List. When the value is 1, the returned data includes the Primary Defect List.
GList	Grown Defect List. When the value is 0, the returned data does not include the Grown Defect List. When the value is 1, the returned data includes the Grown Defect List.
Defect List Format	This field indicates the format of the descriptors returned. If this differs from the format requested in the Command Descriptor Block, a recovered error is returned.
Defect List Length	This indicates the actual number of bytes that follow the defect list header. The drive truncates the list if the allocation length is not sufficient to accommodate the entire list (CHECK CONDITION status is not returned when this happens). The initiator is responsible for comparing the defect list length and the allocation length to ensure that a partial list was not received.

5.26 READ LONG Command (3Eh)

The READ LONG command is used in conjunction with the WRITE LONG command to allow the initiator to issue raw reads and writes of physical data blocks, but address them as logical blocks with an assumed sector length of 512 bytes. These two commands are used to verify the correction capability of the drive's ECC and to create a known kind and location to verify the drive's recovery and bad sector handling policies.

Bit	7	6	5	4	3	2	1	0
Byte								
0	Operation Code (3Eh)							
1	Reserved						Correct	RelAdr
2 - 5	Logical Block Address							
6	Reserved							
7 - 8	Byte Transfer Length (562)							
9	Control							

Figure 5-69 READ LONG Command Descriptor Block — Data Format

Table 5-71 READ LONG Command Descriptor Block — Field Descriptions

Field	Description
Correct	When the value in this field is 1, the drive corrects the data by the ECC before transferring it to the initiator. When the value is 0, the drive does not perform ECC correction on the data. NOTE: This bit, not the DCR bit of the Read Recovery page, governs the effect of error correction for the READ LONG command.
RelAdr	Relative addressing is not supported by the drive. This value must be 0.
Logical Block Address	Specifies the logical sector to read.
Byte Transfer Length	Must be 562. This value reflects the number of bytes spanned by the ECC, including the ECC itself, rounded up to the next word (this is easiest understood in bits: 4096 data bits + 1 force error bit + 15 Logical Block Address bits + 16 cross check bits + 2 pad bits + 360 ECC bits + 6 pad bits that are not written to media but are only needed for transfer = 4496 bits or 562 bytes). Any other value in this field results in a CHECK CONDITION status and the appropriate sense data.

READ LONG Command (3Eh) (continued)

Figure 5-70 shows the format of the data that is returned during the READ LONG Command's data-in buffer transfer. The returned data always consists of 512 bytes of data, followed by a two LBA bytes, two EDC bytes, and 46 ECC-related bytes.

Bit Byte	7	6	5	4	3	2	1	0
0 - 511	Data (512 Bytes)							
512 - 513	Logical Block Address Bits + Force Error Bit							
514 - 515	EDC (8 Bits)							
516 - 562	45 ECC Bits + 8 Pad Bits							

Figure 5-70 READ LONG Command — Returned Data

5.27 READ SKIP MASK Command (E8h)

The READ SKIP MASK command allows the initiator to condition the function of the READ commands to allow only selected data to be transferred over the SCSI interface. A READ SKIP MASK command precedes a READ (10) command; the READ (10) command must be linked to the READ SKIP MASK command or the skip mask is lost and a CHECK CONDITION status is returned.

Bit	7	6	5	4	3	2	1	0
Byte								
0	Operation Code (E8h)							
1	Reserved							
2 - 5	(MSB)	Logical Block Address						(LSB)
6	Skip Mask Length							
7 - 8	(MSB)	Transfer Length						(LSB)
9	VU	Reserved					Flag	Link

Figure 5-71 READ SKIP MASK Command — Data Format

Table 5-72 READ SKIP MASK Command — Field Descriptions

Field	Description
Logical Block Address	This field specifies the first logical block of the corresponding READ (10) COMMAND. This allows a SEEK to begin immediately. The LBA of the following READ (10) command must be the same as the LBA of the READ SKIP MASK command.
Skip Mask Length	This field indicates the length in bytes of the skip mask. The mask is up to 2048 bits, giving a data address capability of 1 MB in 512-byte block length. The mask consists of a sequence of bits where a "1" bit is for the blocks of data to be transferred, and a "0" for the blocks of data that are to be skipped. The first "1" bit of the mask corresponds to the first LBA to be transferred. Any unused bits at the end of the last byte in the mask must be set to 0. The mask is transferred in the data-out buffer transfer of the skip mask command.
Transfer Length	This field must be equal to the transfer length in the linked READ (10) command. The transfer length field must be equal to the number of "1" bits in the skip mask. If a READ (10) is used with 0 as the transfer length, a transfer length of 256 data blocks is implied.

5.28 REASSIGN BLOCKS Command (07h)

The REASSIGN BLOCKS Command reassigns one or more logical blocks to a spare location(s) set aside for this purpose and records the reassignments in the Grown Defect List.

No command-specific Control Descriptor Block fields affect the processing performed for this command.

The list of Logical Block Addresses to reassign is supplied in the REASSIGN BLOCKS Defect List passed in the command's data-out buffer transfer. The Defect List consists of a Defect List Header (Figure 5-73) followed by zero or more Defect Descriptors (Figure 5-74). Table 5-73 contains the field description for the Defect List Header. If multiple Logical Block Addresses are specified, they are listed in ascending order.

The REASSIGN BLOCKS Command Descriptor Block is shown in Figure 5-72.

Bit	7	6	5	4	3	2	1	0
Byte								
0	Operation Code (07h)							
1 - 4	Reserved							
5	Control							

Figure 5-72 REASSIGN BLOCKS Command Descriptor Block — Data Format

Bit	7	6	5	4	3	2	1	0
Byte								
0 - 1	Reserved							
2 - 3	Defect List Length							

Figure 5-73 REASSIGN BLOCKS Defect List Header — Data Format

Table 5-73 REASSIGN BLOCKS Defect List Header — Field Description

Field	Description
Defect List Length	The contents of this field indirectly indicates the number of defective Logical Block Addresses passed in the Defect List. The value must be four times the number of Defect Descriptors supplied. Zero (0) is a valid value and makes the command a NOP.

REASSIGN BLOCKS Command (07h) (continued)

Bit	7	6	5	4	3	2	1	0
Byte								
0 - 3	Defect Logical Block Address							

Figure 5-74 REASSIGN BLOCKS Defect Descriptor — Data Format

The drive returns a CHECK CONDITION status if it cannot complete the command. If this happens because the drive ran out of spare sectors, the drive sets the sense key to MEDIUM ERROR and the additional sense code to NO DEFECT SPARE LOCATION AVAILABLE. The sense key and additional sense key are set appropriately for any other cause. The Logical Block Address of the first defect descriptor not reassigned is returned in the command-specific information field of the sense data. If all defects were reassigned, but some other error occurred, then FFFFFFFFh is posted in the command-specific field.

5.29 RECEIVE DIAGNOSTIC RESULTS Command (1Ch)

The RECEIVE DIAGNOSTIC RESULTS command fetches the results of the last SEND DIAGNOSTIC command sent to the drive. The drives support the diagnostic pages listed in Table 5-74.

Table 5-74 Diagnostic Pages Supported by The Drives

Page Code	Description	Size In Bytes
00h	Supported Diagnostic Pages	6
40h	Translate Address Page (SEND DIAG)	14
40h	Translate Address Page (RECEIVE DIAG)	14 or more

The RECEIVE DIAGNOSTIC RESULTS Command Descriptor Block is shown in Figure 5-75; the data fields are described in Table 5-75.

To receive one or both of the two diagnostic pages that are available, the initiator first sends the desired page code(s) to the drive using a SEND DIAGNOSTICS command. The SEND DIAGNOSTICS command tells the drive which page it should return in response to the next RECEIVE DIAGNOSTICS RESULTS command. The page specified in the most recent SEND DIAGNOSTICS command is always the page returned in response to the RECEIVE DIAGNOSTICS RESULTS command.

Bit	7	6	5	4	3	2	1	0
Byte								
0	Operation Code (1Ch)							
1	Reserved							PCV
2	Page Code Value							
3 - 4	Allocation Length							
5	Reserved			Control				

*Figure 5-75 RECEIVE DIAGNOSTIC RESULTS Command Descriptor Block
— Data Format*

RECEIVE DIAGNOSTIC RESULTS Command (1Ch)

Table 5-75 RECEIVE DIAGNOSTIC RESULTS Command Descriptor Block— Field Descriptions

Field	Description
Allocation Length	Specifies the number of bytes of diagnostic-page results the drive is allowed to send to the initiator. A CHECK CONDITION is not returned if this value truncates the information available.

*RECEIVE DIAGNOSTIC RESULTS Command (1Ch) (continued)***5.29.1 Supported Diagnostics Pages Page (00h)**

Figure 5-76 shows the format of the Supported Diagnostics Pages page.

Bit	7	6	5	4	3	2	1	0
Byte								
0	Page Code (00h)							
1	Reserved							
2 - 3	Page Length (0002h)							
4	Supported Diagnostics Pages Page (00h)							
5	Translate Address Page (40h)							

Figure 5-76 Supported Diagnostics Pages Page — Data Format

To obtain the Supported Diagnostics Pages list, the initiator must:

1. Send a Supported Diagnostics Page to the drive with a SEND DIAGNOSTICS command (the Page Length in the Supported Diagnostics page must be 0).
2. Send a RECEIVE DIAGNOSTIC RESULTS command to the drive. The listing of Supported Diagnostic Pages is returned by the drive.

RECEIVE DIAGNOSTIC RESULTS Command (1Ch) (continued)

5.29.2 Translate Address Page (40h)

The Translate Address Page allows the initiator to translate a logical sector address, physical address, or physical bytes for an index address into any one of the other formats. The address to be translated is passed to the drive with the SEND DIAGNOSTIC command, and the results are returned to the initiator by the RECEIVE DIAGNOSTIC RESULTS command.

Figure 5-77 shows the format of the Translate Address Page that is returned in response to a RECEIVE DIAGNOSTIC RESULTS command, and Table 5-76 contains a description of the data fields. Multiple translated addresses are returned if the logical sector size is greater than 512 bytes (more than one physical data sector).

RECEIVE DIAGNOSTIC RESULTS Command (1Ch) (continued)

Bit Byte	7	6	5	4	3	2	1	0
0	Operation Code (40h)							
1	Reserved							
2 - 3	Page Length (2 + (8 x n)) (where n = number of sectors / LBA)							
4	Reserved					Supplied Format = 000b (LBA)		
5	RAREA	ALTSEC	ALTRK	Reserved		Translate Format = 101b (Physical Sector)		
6 - 8	Translated Address (Cylinder Number) ¹							
9	Translated Address (Head Number) ¹							
10 - 13	Translated Address (Sector Number or Bytes From Index) ¹							
•	•							
•	•							
•	•							
y - y + 2	Translated Address (Cylinder Number) Where y = ((n-1) x 8) + 6							
y + 3	Translated Address (Head Number)							
y + 4 - y + 7	Translated Address (Sector Number or Bytes From Index)							

NOTE: Bytes 6 through 13: One set of these bytes is required for each sector per Logical Block Address.

Figure 5-77 Translate Address Page — Data Format

RECEIVE DIAGNOSTIC RESULTS Command (1Ch) (continued)

Table 5-76 Translate Address Page — Field Descriptions

Field	Description
Supplied Format	The Supplied Format field specifies the format of the following Translate Address field. Must be 000b, which specifies the LBA format.
RAREA	Reserved Area is not used. Must be 0.
ALTSEC	An Alternate Sector bit of 1 indicates that the translated address from the SEND DIAGNOSTIC command is physically located in an alternate sector on the disk. When ALTSEC is 0, it indicates either that no part of the translated address is located in an alternate sector of the medium or that the drive is unable to determine this information.
ALTRK	An Alternate Track bit of 1 indicates either that part or all of the translated address is located on an alternate track of the medium or, the drive cannot determine if all or part of the translated address is located on an alternate track. When ALTRK is 0, it indicates that no part of the translated address is located on an alternate track of the medium
Translate Format	The Translate Format field specifies the format into which the initiator wants the address translated. The value must be 101b, specifying the Physical Sector format.

5.30 RELEASE (6) Command (17h)

The RELEASE and the RESERVE commands are used for contention resolution in multiple-initiator systems. The RELEASE command is used to release a previously reserved logical unit. The drive will not return an error if the initiator attempts to release a reservation that is not currently valid.

Bit	7	6	5	4	3	2	1	0
Byte								
0	Operation Code (17h)							
1	Reserved			3rdPty	Third Party Device ID			Extent
2	Reservation Identification							
3 - 4	Reserved							
5	Control							

Figure 5-78 RELEASE (6) Command Descriptor Block — Data Format

Table 5-77 RELEASE (6) Command — Field Descriptions

Field	Description
3rdPty	Third Party Release allows an initiator to release a logical unit that was previously reserved. If the value in this field is 0, third party release is not requested. When the value is 1, the drive releases the specified logical unit, but only if the reservation was made by the initiator using a third-party reservation that requested the release for the same SCSI device as specified in the Third Party Device ID field.
Third Party Device ID	Required if the 3rdPty bit is 1. This field specifies the SCSI ID (0 through 7 only) of the initiator whose third party reservation is being released. (This field must be set if the initiator of the original third party RESERVE is the source of the RELEASE.)
Extent	The drive supports reservations only on entire logical units. The value must be 0.
Reservation Identification	Any value in this field is ignored by the drive.

NOTE: Assuming that the RELEASE Command Descriptor Block is valid, the drive always returns a GOOD status for this command. An actual release only happens if the initiator has the unit reserved for itself or a third-party initiator.

5.31 RELEASE (10) Command (57h)

The RELEASE and the RESERVE commands are used for contention resolution in multiple-initiator systems. The RELEASE (10) command is used to release a previously reserved logical unit. The drive will not return an error if the initiator attempts to release a reservation that is not currently valid.

Bit	7	6	5	4	3	2	1	0
Byte								
0	Operation Code (57h)							
1	Reserved			3rdPty	Reserved		LongID	Extent
2	Reservation Identification							
3	Third Party Device ID							
4 - 6	Reserved							
7 - 8	(MSB) Parameter List Length							(LSB)
9	Control							

Figure 5-79 RELEASE (10) Command Descriptor Block — Data Format

RELEASE (10) Command (57h) (continued)

Table 5-78 RELEASE (10) Command — Field Descriptions

Field	Description
3rdPty	Third Party Release allows an initiator to release a logical unit that was previously reserved. If the value in this field is 0, third party release is not requested. If 3rdPty = 1, then the device server shall release the specified logical unit, but only if the initiator ID, 3rdPty bit, and third party device ID are identical when compared to the RESERVE command that established the reservation.
LongID	If the Long ID bit is set to 1, the Parameter List Length is 8 and the eight bytes of the parameter list carry the device ID of the third party device; the contents of the Third Party Device ID in the CDB (byte 3) are ignored.
Third Party Device ID	If the Third Party Device ID value that is connected with the reservation release is smaller than 255, the LongID bit may be 0 and the ID value sent in the CDB. If LongID bit = 0, the Parameter List Length field also = 0. If the Third Party Device ID value is greater than 255, LongID = 1. Device servers that support device IDs greater than 255 will accept commands with LongID = 1; device servers that do not support IDs greater than 255 may reject commands with LongID set = 1. Device ID formats are protocol-specific.
Parameter List Length	The contents of this field specify the length, in bytes, of the parameter list that will be transferred from the initiator to the target.
Extent	The drive supports reservations only on entire logical units. The value must be 0.
Reservation Identification	Any value in this field is ignored by the drive.

NOTE: Assuming that the RELEASE Command Descriptor Block is valid, the drive always returns a GOOD status for this command. An actual release only happens if the initiator has the unit reserved for itself or a third-party initiator.

5.32 REPORT DEVICE IDENTIFIER Command (A3h)

The REPORT DEVICE IDENTIFIER command requests that the device server send device identification to the application client

Bit	7	6	5	4	3	2	1	0
Byte								
0	Operation Code (A3h)							
1	Reserved			Service Action (05h)				
2 – 5	Reserved							
6 – 9	Allocation Length (MSB) (LSB)							
10	Reserved							
11	Control							

Figure 5–80 REPORT DEVICE IDENTIFIER Command Descriptor Block — Data Format

Table 5–79 REPORT DEVICE IDENTIFIER Command Descriptor Block—Field Descriptions

Field	Description
Service Action	Must = 05h. Any other value forces Check Condition, Illegal Request.
Allocation Length	If the length is not sufficient to contain all of the parameter data, the first portion of the data is returned, and this is not considered to be an error. The actual length of the parameter data is available in the Identifier Length field in the parameter data. If the remainder of the parameter data is required, the application client must send a new REPORT DEVICE IDENTIFIER command with an Allocation Length field large enough to contain all of the data.

REPORT DEVICE IDENTIFIER Command (A3h) (continued)

The REPORT DEVICE IDENTIFIER command's parameter list contains a 4-byte field that provides the length, in bytes, of the parameter list and the logical unit's identifier.

Bit	7	6	5	4	3	2	1	0
Byte								
0 - 3	(MSB) Identifier Length (n - 4)							(LSB)
4 - n	Identifier							

Figure 5-81 REPORT DEVICE IDENTIFIER Parameter List — Data Format

Table 5-80 REPORT DEVICE IDENTIFIER Parameter List — Field Descriptions

Field	Description
Identifier Length	This field specifies the length, in bytes, of the Identifier field. If the Allocation Length field in the CDB is insufficient in size to transfer all of the Identifier, the length is not adjusted to reflect any truncation of information. The Identifier is initially set = 0, and is changed only by a successful SET DEVICE IDENTIFIER command.
Identifier	The value reported in this field is a vendor-specific value and is the last value written by a successful SET DEVICE IDENTIFIER command. The value of the Identifier is changed only by a successful SET DEVICE IDENTIFIER command. The Identifier value persists through device resets, power cycles, and formatting of media.

5.33 REPORT LUNS Command (A0h)

The REPORT LUNS command requests that the peripheral device logical unit numbers of known logical units in the target disk be sent to the applications client. The command only returns information about the logical units to which commands may be sent.

Bit	7	6	5	4	3	2	1	0
Byte								
0	Operation Code (A0h)							
1 - 5	Reserved							
6 - 9	(MSB) Allocation Length (LSB)							
10	Reserved							
11	Control							

Figure 5-82 REPORT LUNS Command Descriptor Block — Data Format

Table 5-81 REPORT LUNS Command Descriptor Block—Field Description

Field	Description
Allocation Length	<p>If the Allocation Length is not sufficient to contain the logical unit number values for all configured logical units, the device server still reports as many logical number values as will fit in the Allocation Length.</p> <p>The format of the report of configured logical units is shown in Figure 5-83.</p>

REPORT LUNS Command (0Ah) (continued)

Bit	7	6	5	4	3	2	1	0
Byte								
0 - 3	(MSB) LUN List Length (n - 7)							(LSB)
4 - 7	Reserved							
8 - 15	(MSB) LUN (first LUN)							(LSB)
•								•
•								•
•								•
n-7 - n	LUN (last LUN, if more than one)							

Figure 5-83 LUN Reporting Parameter List — Data Format

The LUN List Length field contains the length in bytes of the LUN list that can be transferred. The LUN list length equals the number of logical unit numbers reported multiplied by eight. If the allocation length in the CDB is too small to allow transfer of information about all of the logical units configured, the LUN list length value is not adjusted to reflect the truncation.

5.34 REQUEST SENSE Command (03h)

The REQUEST SENSE command causes the drive to transfer detailed sense data to the initiator. The drive maintains sense data on an individual initiator basis.

Bit	7	6	5	4	3	2	1	0
Byte								
0	Operation Code (03h)							
1 - 3	Reserved							
4	Allocation Length							
5	Control							

Figure 5-84 REQUEST SENSE Command Descriptor Block — Data Format

Table 5-82 REQUEST SENSE Command Descriptor Block—Field Description

Field	Description
Allocation Length	The drives can return a maximum of 18 bytes of sense data. The initiator should set this field to the maximum value to receive all the sense data.

5.34.1 Sense Data Availability

Sense data is available if the initiator is responsible for an in-progress FORMAT UNIT command (Format Progress Indication is available in the Format Progress Indication Bytes).

The drive returns a sense key of NO SENSE and an additional sense code of NO ADDITIONAL SENSE INFORMATION if it has no sense data available for the initiator.

For Parallel SCSI disk drives, sense data is available if the previous command from the initiator terminated with a CHECK CONDITION status, or if the previous command from the initiator ended with an unexpected BUS FREE error.

REQUEST SENSE Command (03h) (continued)

5.34.2 Clearing Sense Data

Sense data is cleared from the drive's memory by one of the following conditions:

- After being returned from a REQUEST SENSE command.
- As soon as any other command is received from the initiator except INQUIRY.

5.34.3 Status Reporting

The drive returns CHECK CONDITION status for a REQUEST SENSE command only to report errors specific to the command when:

- A non-zero reserved bit is detected in the Command Descriptor Block.
- A drive malfunction prevents return of the sense data.

For Parallel SCSI disk drives, the drive also returns CHECK CONDITION status for a REQUEST SENSE command to report errors specific to the command when an unrecovered parity error is detected on the data bus.

Note that any of the previous conditions overwrites the current sense data with the sense data describing the error encountered by the REQUEST SENSE command (that is, for the condition stated in the first bulleted entry, the new sense data reflects that a non-zero reserved bit was detected in the Command Descriptor Block).

The drive returns a maximum of 18 bytes of sense data. The initiator should set the Allocation Length field in the Command Descriptor Block to 18 to receive all sense data. If the Allocation Length is set to 0, four bytes of data will be returned (this option maintains compatibility with SCSI-1 systems).

5.34.4 Sense Data Format for Error Code 70h and Error Code 71h

The sense data format for error code 70h (current error) and error code 71h (deferred error) is shown in Figure 5-85; the data fields are defined in Table 5-83.

REQUEST SENSE Command (03h) (continued)

Bit Byte	7	6	5	4	3	2	1	0
0	Valid	Error Code (70h or 71h)						
1	Segment Number = 00h							
2	FileMrk = 0	EOM = 0	ILI	Rsvd	Sense Key			
3 - 6	Information							
7	Additional Sense Length = 0Ah							
8 - 11	Command-Specific Information							
12	Additional Sense Code							
13	Additional Sense Code Qualifier							
14	Field Replaceable Unit Code							
15	SKSV	Sense-Key Specific						
16 - 17	Sense-Key Specific (continued)							

Figure 5-85 Sense Data Format for Error Code 70h or 71h — Data Format

REQUEST SENSE Command (03h) (continued)

Table 5–83 Sense Data Fields (Error Code 70h) — Field Descriptions

Field	Definition
Valid	The Valid bit, when set to 1, indicates that the Information field contains valid information, as defined below. When set to 0, it indicates that the contents of the Information field are not as defined below.
Error Code	Either 70h, indicating that current errors are to be reported, or 71h, indicating deferred errors are to be reported.
Segment Number	This feature is not supported and will return a zero value.
FileMrk	The File Mark field is reserved for sequential access devices. This field will return a zero value.
EOM	The End-of-Medium field is reserved for sequential access and printer devices. This field will return a zero value.
ILI	Incorrect Length Indicator. When this bit is set to 1, it indicates that the requested logical sector length did not match the logical sector length of the data on the media. This field is only set for incorrect byte transfer lengths on a READ LONG and WRITE LONG command descriptor block.
Sense Key	The sense key, additional sense code, and additional sense code qualifier provide a hierarchy of information. The sense key provides the highest level of information about error and exception conditions. The additional sense code and qualifier provided more detailed information. The sense keys supported by the drive are described in Table 5–84.
Information	The Information field contains the following device-type or command-specific information: <ul style="list-style-type: none"> 1) The unsigned logical sector address associated with the sense key. 2) The difference between the requested length and the actual length in either bytes or sectors, as determined by the command. This difference is called the residue. 3) A description of each of the drive's supported commands definition for this field is contained in Table 5–85.
Additional Sense Length	The Additional Sense Length field indicates the number of additional sense bytes that can follow. If the Allocation Block Length specified in the Command Descriptor Block is too small to transfer all of the additional sense bytes, the Additional Sense Length byte value is not adjusted to reflect the truncation.
Command-Specific Information	The Command-Specific Information field is valid only for Sense Data for a failed REASSIGN BLOCKS command. It contains either the Logical Block Address of the first defect descriptor not reassigned or, if all the defect descriptors were reassigned, FFFFFFFh.
Additional Sense Code	The Additional Sense Code field indicates further information related to the error or exception condition reported in the sense key. If the drive does not have further related information, the additional sense key is set to NO ADDITIONAL SENSE INFORMATION. See Table 5–86.
Additional Sense Code Qualifier	The Additional Sense Code Qualifier field indicates detailed information related to the additional sense code. If the drive does not have detailed information, the field is set to 0.

(continued)

REQUEST SENSE Command (03h) (continued)

Table 5-83 Sense Data Fields (Error Code 70h) — Field Descriptions (continued)

Field	Definition								
Field Replaceable Unit	The Field Replaceable Unit Code specifies a unit that has failed: <table border="1"> <thead> <tr> <th>Code</th> <th>FRU</th> </tr> </thead> <tbody> <tr> <td>1</td> <td>Host</td> </tr> <tr> <td>2</td> <td>Hard Disk Assembly (HDA)</td> </tr> <tr> <td>3</td> <td>Printed Circuit Board Assmby (PCBA)</td> </tr> </tbody> </table>	Code	FRU	1	Host	2	Hard Disk Assembly (HDA)	3	Printed Circuit Board Assmby (PCBA)
Code	FRU								
1	Host								
2	Hard Disk Assembly (HDA)								
3	Printed Circuit Board Assmby (PCBA)								
SKSV	Sense-Key Specific Valid bit. This bit is 1 for all sense keys listed in Table 5-84 (exception: NOT READY and UNIT ATTENTION/RPL Status Change). The SKSV bit and Sense-Key Specific bytes may be 0 when the sense key is NOT READY.								
Sense-Key Specific	The contents of the Sense-Key Specific field are unique to each sense key. Table 5-85 summarizes the sense keys for which the drive supplies sense-key specific information. (See SKSV.)								

Table 5-84 Supported Sense Keys

Sense Key	Code	Description
NO SENSE	0h	There is no specific sense key information to be reported for the designated logical unit. This is the case for a successful command or a command that received CHECK CONDITION or COMMAND TERMINATED status because one of the following bits is set to 1: File Mark, EOM, or ILI.
RECOVERED ERROR	1h	The last command completed successfully with some recovery action performed by the drive. Details may be contained in the additional sense bytes and the Information field. When multiple recovered errors occur during one command, the choice of which error to report (first, last, most severe, etc.) is device specific.
NOT READY	2h	The logical unit cannot be accessed. Operator intervention may be required to correct this condition.
MEDIUM ERROR	3h	The command terminated with a non-recovered error condition that was probably caused by a flaw in the medium or an error in the recorded data. This sense key can also be returned if the drive is unable to distinguish between a flaw in the medium and a specific hardware failure (Sense Key 4h).
HARDWARE ERROR	4h	The drive detected a non-recoverable hardware failure while performing the command or during a self-test. The error type could be controller failure, device failure, etc.

(continued)

REQUEST SENSE Command (03h) (continued)

Table 5-84 Supported Sense Keys (continued)

Sense Key	Code	Description
ILLEGAL REQUEST	5h	The drive detected an illegal parameter in the Command Descriptor Block or in the additional parameters supplied as data for some commands. If the drive detects an invalid parameter in the Command Descriptor Block, it terminates the program without altering the medium. If the drive detects an invalid parameter in the additional parameters supplied as data, the drive may have already altered the medium. This sense key may also indicate that an invalid IDENTITY message was received.
UNIT ATTENTION	6h	The drive generates this sense key whenever the drive is reset by a TARGET RESET message, a hard reset condition, or by a power-on reset. Other conditions that generate a UNIT ATTENTION include: <ul style="list-style-type: none"> • The mode parameters in effect for this initiator have been changed by another initiator. • The microcode has been changed. • Tagged commands queued for this initiator were cleared by another initiator. • INQUIRY data has been changed. • The mode parameters in effect for the initiator have been restored from non-volatile memory. • Other events that require the attention of the initiator.
DATA PROTECT	7h	A Read or Write command was attempted on a sector that is protected for this operation. The requested operation was not performed.
VENDOR SPECIFIC	9h	This sense key is used to report specific Quantum conditions.
ABORTED COMMAND	Bh	The drive aborted the command. The initiator may be able to reissue the command successfully.
MISCOMPARE	Eh	The source data did not match the data read from the medium (i.e., a mismatch).

REQUEST SENSE Command (03h) (continued)

Table 5-85 Sense Key Information Field Contents

Command	Sense Key	Field Contents
FORMAT UNIT	RECOVERED ERROR	The address of the first bad sector encountered during the format operation. This sector, and all other bad sectors found during a format operation, are automatically reallocated by the drive.
	ILLEGAL REQUEST	The first illegal defect descriptor encountered in the Defect Descriptor List.
READ(6)(10)	RECOVERED/MEDIUM ERROR	The address of the first bad sector found during the read (possible only when PER=1, DTE=0, and DCR=0.). For a RECOVERED ERROR, the last Logical Block Address in the error is returned; for a MEDIUM ERROR the first Logical Block Address is returned.
READ LONG	RECOVERED/MEDIUM ERROR	READ(6)(10) error recovery qualifiers apply. The information field echoes this command's Logical Block Address field.
	ILLEGAL REQUEST	The transfer length residue (requested, minus correct length) if the transfer length is wrong. Note: ILI will be set.
REASSIGN BLOCKS	HARDWARE ERROR	The address of the first logical sector is not reassigned (the drive ran out of spare sectors).
	ILLEGAL REQUEST	The first illegal defect descriptor encountered in the Defect Descriptor List.
SEEK(6)(10)	ILLEGAL REQUEST	The first invalid sector address detected by the command.
VERIFY	RECOVERED/MEDIUM ERROR	WRITE(6)(10) error recovery qualifiers apply. Contains the Logical Block Address that may have failed the write or verify (compare) operation. For a RECOVERED ERROR, the last Logical Block Address in the error is returned; for a MEDIUM ERROR the first Logical Block Address is returned.

(continued)

REQUEST SENSE Command (03h) (continued)

Table 5-85 Sense Key Information Field Contents (continued)

Command	Sense Key	Field Contents
WRITE(6)(10)	RECOVERED/MEDIUM ERROR	The address of the first bad sector found during the write. (Possible only when PER=1, DTE=0, and DCR=0.) For a RECOVERED ERROR, the last Logical Block Address in the error is returned; for a MEDIUM ERROR the first Logical Block Address is returned.
WRITE AND VERIFY	RECOVERED/MEDIUM ERROR	WRITE(6)(10) error recovery qualifiers apply. This field contains the Logical Block Address that may have failed the write or verify (compare) operation. For a RECOVERED ERROR, the last Logical Block Address in the error is returned; for a MEDIUM ERROR the first Logical Block Address is returned.
WRITE LONG	RECOVERED/MEDIUM ERROR	Write (10) error recovery qualifiers apply. The information field echoes the WRITE LONG command's Logical Block Address.
	ILLEGAL REQUEST	The transfer length residue (requested minus correct length) if the transfer length is wrong. Note: ILI will be set.

REQUEST SENSE Command (03h) (continued)

Table 5-86 Supported Additional Sense Codes and Sense Code Qualifiers (in Hex)

Sense Key	Sense Code	Sense Qualifier	Meaning	
00h NO SENSE	00	00	No Specific Sense Key Information to be Reported for Designated Logical Unit	
	0B	01	SMART Overtemperature Trip	
	5D	00	SMART Tripped	
	5D	01	SMART Triggered due to READ or WRITE TA Factor	
	5D	02	SMART Triggered due to READ or WRITE Error Factor	
	5D	0B	Spin Failure Retry Threshold Exceeded	
	5D	FF	SMART Test Trip	
	5E	00	Power Management Enabled	
	5E	01	Power Management Enabled, Drive is in IDLE Sub-State (Idle Timer Timed Out)	
	5E	02	Power Management Enabled, Drive is in STANDBY Sub-State (Standby Timer Timed Out)	
	5E	03	Power Management State = "Sleep"	
	01h RECOVERED ERROR	02	00	No Seek Complete
		03	00	Write Fault, Recovered
03		01	Pre-Amp WRITE Unsafe, Recovered	
09		00	Track Following Error, Recovered	
0B		01	SMART Overtemperature Trip	
0C		00	Write Error, Recovered	
0C		01	Write Error, Recovered with Auto Reallocation	
15		02	Positioning Error Detected by Read of Medium	
15		03	Positioning Error Detected by Read, Recovered	
15		04	Servo Write Not OK Assertion, Recovered	
15		05	Servo Read Not OK Assertion, Recovered	
15		06	Spoke Address Error, Recovered	
15		80	Positioning Error Detected by READ on WRITE: Recovered	
15		81	Servo WRITE Not OK Assertion: Recovered	
15		82	Servo READ Not OK Assertion: Recovered	
16	00	Data Synchronization Mark Error		
16	05	Spoke Synchronization Mark Error, Recovered		
16	83	Spoke Synchronization Mark Error, Recovered		

(continued)

REQUEST SENSE Command (03h) (continued)

Table 5-86 Supported Additional Sense Codes and Sense Code Qualifiers (in Hex)
(continued)

Sense Key	Sense Code	Sense Qualifier	Meaning
01h RECOVERED ERROR (continued)	17	00	Recovered Data, No Error Correction Applied
	17	01	Recovered Data With Retries
	17	AF	Recovered Data with ECC On-the-Fly
	18	00	ECC Error
	18	02	Recovered Data, Data Auto Reallocated
	18	05	Recovered Data, Recommend Reassignment
	1C	00	Defect List Not Found
	40	86	Controller FIFO Over/Underflow Error, Recovered
	40	8C	LBA Mismatch Error, Recovered
	41	86	Required H/W Process Did Not Complete, Recovered
	42	00	Power On or Self Test Failed
	5D	00	SMART Tripped
	5D	01	TA Error Rate Threshold Exceeded
	5D	02	READ or WRITE Error Rate Thresh Exceeded
	5D	0B	Spin Failure Retry Threshold Exceeded
	5D	FF	Test Trip
	80	10	Block Slow PLL Error, Recovered
	80	11	Split Slow PLL Error, Recovered
	80	12	Block State Machine Disabled, Recovered
	82	20	Motor Speed Error, Recovered
02h NOT READY	04	00	Drive Not Ready
	04	01	Logical Unit is in the Process of Becoming Ready
	04	02	Logical Unit is Not Ready, Initializing Command Required
	04	04	Logical Unit Is Not Ready, Format In Progress
	04	80	FileSystem Boot Area Unreadable at POR
	04	81	Drive Not Ready, Servo Error (See Info Block)
	04	82	Failure to Read Sector Descriptor File

(continued)

REQUEST SENSE Command (03h) (continued)

Table 5-86 Supported Additional Sense Codes and Sense Code Qualifiers (in Hex)
(continued)

Sense Key	Sense Code	Sense Qualifier	Meaning
02h NOT READY (continued)	04	83	Error Reading Diskware File
	04	84	Diskware Not Loaded Yet – Issue Start CMD
	04	85	Code Jumped to Dead Vector
	04	86	Drive Not Ready, Configuration Page 8/HDA Mismatched
	04	88	Error Reading Config Page File
	04	91	Drive Not Ready, Self Scan is Running
	19	00	Defect List in Error
	19	01	Drive Not Ready, Unknown Defect List
	19	03	Defect List Error in Grown List
	31	00	Drive Not Ready, Format Corrupted
	31	01	FORMAT UNIT Command Failed
	40	80	Diagnostics Failure, Unknown Reason
	40	81	Diagnostics Failure, RAM Failure
	40	82	Diagnostics Failure, Inconsistent Data
	40	83	Diagnostics Failure, HDA Diagnostic
	40	84	Diagnostics Failure, PCB Diagnostic
	40	85	Error Detection Code (EDC) Error
	40	86	DMA Overrun Error
	40	99	ROM and Resident Code are Incompatible
	03h MEDIA ERROR	02	00
03		01	Pre-Amp Write Unsafe
03		84	Diag: Write Fault
0C		00	Write Error
0C		02	Write Error, Auto Reallocation Failed
0C		03	Write Error, Recommend Reassignment
11		00	Unrecovered Read Error
11		04	Unrecovered Read Error, Auto Reallocation Failed
11		0B	Unrecovered Read Error, Recommend Reassignment
11		81	System File Has Checksum Error
	11	84	Error Reading Diskware Vector Table File

(continued)

REQUEST SENSE Command (03h) (continued)

Table 5-86 Supported Additional Sense Codes and Sense Code Qualifiers (continued)

Sense Key	Sense Code	Sense Qualifier	Meaning
03h MEDIA ERROR (continued)	15	00	Positioning Error Detected by Drive, Check UPT
	15	02	Positioning Error Detected by Read of Medium
	15	80	Positioning Error Detected by Write of Medium
	15	81	Servo Write Not OK
	15	82	Servo Read Not OK
	15	83	Spoke Address Error
	16	00	Data Synchronization Mark Error
	16	83	Spoke Synchronization Mark Error
	18	00	Recoverable ECC Error
	19	00	Defect List in Error
	19	01	Defect List Error, Unknown Format
	19	03	Defect List Error in Grown List
	19	80	Replacement Block Bad
	19	81	Found Two Bad RBNs
	27	00	Grown List Error or Incomplete Replacement, Unit Write Protected
	30	01	Cannot Read Medium, Unknown Format
	31	00	Medium Format Corrupted
	31	01	FORMAT Command Failure
	32	00	No Defect Spare Location Available
	32	01	Defect List Update Failure
	41	86	Required H/W Process Did Not Complete
	44	80	Error Writing FileSystem (neg cyl) File
	44	81	Error Reading FileSystem (neg cyl) File
	80	00	Forced Error
	80	10	
	80	11	Split Slow PLL Error
	80	12	Block State Machine Disabled Error
82	20	Motor Speed Error	

(continued)

REQUEST SENSE Command (03h) (continued)

Table 5-86 Supported Additional Sense Codes and Sense Code Qualifiers (continued)

Sense Key	Sense Code	Sense Qualifier	Meaning
04h HARDWARE ERROR	00	00	No Additional Sense Information
	02	00	Could Not Complete SEEK
	03	00	Write Fault
	03	84	Diagnostics Failed, Unit is Write Protected
	06	80	PES Calibration Results Are Bad
	06	81	BIAS Calibration Results Are Bad
	06	82	GAMMA Calibration Results Are Bad
	06	83	Recalibration Record Timeout
	06	84	Seek Fail During Feedforward Scal Calibration
	06	85	Initial Track Grab Failed
	06	86	Seek Failure During Bias Calibration
	06	87	Seek Failure During FT/J Calibration
	06	88	Seek Failure During PES Calibration
	06	89	Recalibration Servo Not Ready
	06	8A	Switch Point Calibration Results Are Bad
	06	8B	Recalibration Track Number Bad
	06	8C	Seek Failure During Switch Point Calibration
	06	8D	Not All Head(s) in USER Head Map Were Found Alive
06	8E	Valid SAM(s) Not Found	

(continued)

REQUEST SENSE Command (03h) (continued)

Table 5-86 Supported Additional Sense Codes and Sense Code Qualifiers (continued)

Sense Key	Sense Code	Sense Qualifier	Meaning
04h HARDWARE ERROR (continued)	06	8F	Sync to Disk Spokes Failed
	06	90	Invalid CP18 on Disk
	06	91	Invalid CP21 on Disk
	06	92	Invalid CP22 in NVR
	06	93	Head Polarity Calibration Failed
	06	94	Invalid Head Polarity Vector in CP22
	06	95	Seek Failure During Servo Writer Erase Band Resync
	06	99	Seek 2 Failure During Servo Writer Erase Band Resync
	06	9A	Seek Failure During pes Linearization Table Update
	06	9B	Motor Speed Error After Turbo Spin Recovery
	09	00	Track Following Error
	09	90	Svo, No Error
	09	91	Svo, TA Err – 1 st
	09	92	Svo, TA Err – 2 nd
	09	93	Svo, Bad Sync – 1 st
	09	94	Svo, Bad Sync – 2 nd
	09	95	Svo, Grey Code Quality Error – 1 st
	09	96	Svo, Grey Code Quality Error – 2 nd
	09	97	Svo, Head Spoke Number Miscompare – 1 st
	09	98	Svo, Head Spoke Number Miscompare – 2 nd
	09	A0	Svo, Possible Bump – 1 st
	09	A1	Svo, Possible Bump – 2 nd
	09	A2	Svo, Defective Wdg – 1 st
	09	A3	Svo, Defective Wdg – 2 nd
	09	A4	Svo, Mapped Wdg – 1 st
	09	A5	Svo, Mapped Wdg – 2 nd
	09	A6	Svo, ID Error – 1 st
	09	A7	Svo, ID Error – 2 nd
	09	A8	Svo, Unknown Wdge Err – 1 st
	09	A9	Svo, Unknown Wdge Err – 2 nd
	09	B0	Svo, Timeout Waiting for Burst
	09	B1	Svo, Confined Sync Timeout
	09	B2	Svo, R/W Serial Timeout
	09	B3	Svo, Context Acquire Timeout

(continued)

REQUEST SENSE Command (03h) (continued)

Table 5-86 Supported Additional Sense Codes and Sense Code Qualifiers (continued)

Sense Key	Sense Code	Sense Qualifier	Meaning
04h HARDWARE ERROR (continued)	09	B4	Svo, Context Acquire Data Address Valid Timeout
	09	B5	Svo, Context Acquire Bad Track Number
	15	00	Random Positioning Error
	15	02	Positioning Error Detected By Read
	32	00	No Defect Spare Location Available
	32	80	Replacement Failed (Table May Be Full)
	35	00	ESI - Unspecified Enclosure Services Failure
	35	01	ESI- Unsupported Enclosure Function
	35	02	ESI - Enclosure Services Unavailable
	35	03	ESI - Enclosure Transfer Failure
	35	04	ESI – Enclosure Transfer Refused
	40	80	Diagnostics Failure, Unknown Reason
	40	81	Diagnostics Failure, RAM Failure
	40	82	Diagnostics Failure, Inconsistent Data or Bugcheck Error code in information field.
	40	83	Diagnostics Failure, HDA Diagnostic
	40	84	Diagnostics Failure, PCB Diagnostic
	40	85	Error Detection Code (EDC) Error
	40	86	DMA Failure
	40	87	EDC Check Failed
	40	88	LBA Error
	40	89	Failed XOR Test
	40	8A	Failed Servo Engine Test
	40	8C	LBA Mismatch Error
	40	8D	ETM Test Failure, Drive Failed to Spin
	40	8E	Could Not Achieve Target RPM
	40	A0	Test Did Not Complete
	40	A1	Test Completed But Failed
	40	A2	Invalid Parameter Was Received
	40	A3	Invalid Data in Servo Record Prevented Algorithm From Completing
	40	A4	Servo Record Did Not Complete
	40	A5	Read Seek Failed
	40	A6	Write Seek Failed
	40	A7	Raw Seek Failed
	40	A8	Servo Diagnostic Test Running
40	A9	Invalid 97/98 Super CMD sequence; Must Send 97 Before 98	
40	AA	Super CMD 97 Already Running	

(continued)

REQUEST SENSE Command (03h) (continued)

Table 5-86 Supported Additional Sense Codes and Sense Code Qualifiers (continued)

Sense Key	Code	Qualifier	Meaning
04h HARDWARE ERROR (continued)	40	AB	Could Not Calculate Bode Info
	40	AC	Could Not Complete the Low-Level BCV WRITE Function
	40	AD	Got a Divide By Zero Error in Complex Divide
	41	00	Data Path Failure
	41	86	DDMA Overrun or ACK Overrun
	44	00	Logical Assertion (Firmware Consistency Check)
	44	84	Insufficient Buffer Memory for Operation
05h ILLEGAL REQUEST	1A	00	Parameter List Length Error
	20	00	Invalid Command Operation Code
	20	80	R4K Image Loader Bad
	20	81	R4K System Micro-OCM Code Image (Not in EEROM) Bad
	20	82	R4K Diagnostic Code Image (Not in EEROM) Bad
	20	FE	ETM Cannot Run With Motor Spinning
	20	FF	Selfscan Code Not Loaded or Bad
	21	00	Logical Block Address Out of Range
	24	00	Invalid Field in CDB
	25	00	Logical Unit Not Supported
	26	00	Invalid Fields in Parameters
	26	02	Invalid Parameter Value
	26	04	Invalid Release of Persistent Reservation
	2C	00	Command Sequence Error
	3D	00	Invalid Bits in IDENTIFY Message
49	00	Invalid Message Error	
06h UNIT ATTENTION	0B	01	SMART Overtemperature Trip
	28	00	Format Completed by Another Initiator
	29	00	Power On, SCSI Bus Reset, or Bus Device Reset Occurred
	29	01	Power On Occurred
	29	02	SCSI Bus Reset Occurred

(continued)

REQUEST SENSE Command (03h) (continued)

Table 5-86 Supported Additional Sense Codes and Sense Code Qualifiers (continued)

Sense Key	Code	Qualifier	Meaning	
06h UNIT ATTENTION (continued)	29	03	Bus Device Reset Occurred	
	29	04	Device Reset Occurred, Bugcheck Occurred	
	29	05	Bus Mode Changed to Single-Ended	
	29	06	Bus Mode Changed to LVDS	
	29	86	Commands Aborted, ACK Overrun Detected	
	2A	01	Mode Parameters Changed	
	2A	02	Log Parameters Changed	
	2A	03	Persistent Reservations Were Cleared	
	2F	00	Commands Aborted by Another Initiator	
	3F	01	Target Operating Condition Changed, Microcode has Changed	
	5C	01	Spindle Sync'd	
	5C	02	Spindle Not Sync'd	
	5D	00	SMART Tripped	
	5D	01	TA Error Rate Threshold Exceeded	
	5D	02	READ or WRITE Error Rate Thresh. Exceeded	
	5D	0B	SMART Triggered Due to Spin Retry	
	5D	FF	SMART Test Trip	
	07h DATA PROTECT	27	00	Unit Hardware or Firmware Write Protected
	09h VENDOR UNIQUE	81	00	Read Overlaps Deferred Write
	0Bh ABORTED COMMAND	00	00	No Additional Sense Information Error
25		00	Logical Unit Not Supported	
41		00	Data Path Failure	
41		86	Commands Aborted, ACK Overrun Detected*	
43		00	Message Reject Error	
44		00	Busy Timeout	
45		00	Selection or Reselection Failure	
47		00	SCSI Bus Parity Error Detected	
48		00	Initiator Detected Error (IDE) Message Received	
48		86	ACK Overrun Detected Error	
49		00	Invalid Message Error	
4E	00	Overlapped Commands Attempted		
0Eh MISCOMPARE	1E	00	No Additional Information	

Table 5-87 Sense-Key Specific Field Contents

Sense Key	Field Contents	Initiator Interpretation
ILLEGAL REQUEST	Field pointer bytes (Figure 5-86)	Illegal field's offset (in the Command Descriptor Block or data-out buffer transfer parameters).
NOT READY	Progress Indication (Figure 5-87)	Indicates progress of FORMAT UNIT Command.
RECOVERED ERROR	Retry count (Figure 5-88)	Number of retries for I/O operation
MEDIUM ERROR	Retry count (Figure 5-88)	Number of retries for I/O operation

The data format of the Field Pointer bytes associated with the ILLEGAL REQUEST sense key is shown in Figure 5-86; the description of the data is in Table 5-88.

REQUEST SENSE Command (03h) (continued)

Bit	7	6	5	4	3	2	1	0
Byte								
15	SKSV = 1	C/D	Reserved		BPV	Bit Pointer		
16 - 17	Field Pointer							

Figure 5-86 ILLEGAL REQUEST Sense Key Field Pointer Bytes — Data Format

Table 5-88 ILLEGAL REQUEST Sense Key Field Pointer Bytes — Field Descriptions

Field	Description
SKSV	Sense-Key Specific Valid
C/D	Command/Data. When the value in this field is 1, the Field Pointer identifies the first byte in error in the Command Descriptor Block. When 0, the Field Pointer identifies the first byte in error in the data parameters sent by the initiator in the data-out buffer transfer.
BPV	Bit Pointer Valid. The drive does not support Bit-level resolution; the bit value must be 0.
Bit Pointer	The drive does not support Bit-level resolution; the bit value must be 0.
Field Pointer	Indicates the first byte in error in either the Command Descriptor Block or in the data parameters sent by the initiator.

The data format of the Progress Indication bytes associated with the NOT READY sense key is shown in Figure 5-87; the description of the data is contained in Table 5-89.

Bit	7	6	5	4	3	2	1	0
Byte								
15	SKSV = 1	Reserved						
16 - 17	Progress Indication							

Figure 5-87 NOT READY Sense Key - Progress Indication Bytes — Data Format

REQUEST SENSE Command (03h) (continued)

Table 5-89 NOT READY Sense Key - Progress Indication Bytes — Field Descriptions

Field	Description
SKSV	Sense-Key Specific Valid
Progress Indication	Indication of percent complete for the FORMAT UNIT command. The returned value is the numerator of a fraction that has 65536 (10000h) as the denominator (Progress Indication/10000h = percent complete).

The data format of the Retry Count bytes associated with the MEDIUM ERROR or RECOVERED ERROR sense key is shown in Figure 5-88; the description of the data is contained in Table 5-90.

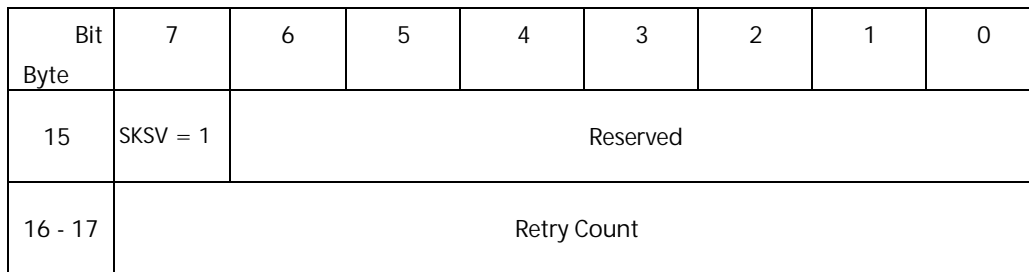


Figure 5-88 MEDIUM ERROR or RECOVERED ERROR Sense Key - Retry Count — Data Format

Table 5-90 MEDIUM ERROR or RECOVERED ERROR Sense Key - Retry Count — Field Descriptions

Field	Description
SKSV	Sense-Key Specific Valid
Retry Count	The number of times an I/O operation was retried

5.35 RESERVE (6) Command (16h)

The RESERVE and the RELEASE commands are used for contention resolution in multiple-initiator systems. The RESERVE command is used to reserve a logical unit.

Bit	7	6	5	4	3	2	1	0
Byte								
0	Operation Code (16h)							
1	Reserved			3rdPty	Third Party Device ID			Extent
2	Reservation Identification							
3 - 4	Extent List Length							
5	Control							

Figure 5-89 RESERVE (6) Command Descriptor Block — Data Format

Table 5-91 RESERVE (6) Command — Field Descriptions

Field	Description
3rd Pty	When set to 1, it indicates that the Third Party Device ID field is valid.
Third Party Device ID	<p>Required and used only when the 3rdPty bit is set, in which case this field specifies the SCSI ID (0 through 7 only) of the initiator to be granted the reservation of the logical unit. The drive preserves the reservation until one of the following occurs:</p> <ul style="list-style-type: none"> • It is superseded by another valid RESERVE command from the initiator. • It is released by the same initiator. • It is released by a TARGET RESET message from any initiator. • It is released by a hard reset condition. <p>The drive ignores any attempt to release the reservation made by any other initiator. For example, if ID7 sends ID2 a Third Party reservation on behalf of ID6 (the target at ID2 gets reserved for the initiator ID6), then only ID7 can release the target at ID2 (using a Third Party release). ID6 cannot release the reservation even though the reservation was made on its behalf.</p>

(continued)

RESERVE (6) Command (16h) (continued)

Table 5-91 RESERVE (6) Command — Field Descriptions (continued)

Field	Description
Extent	The drive supports reservations only on entire logical units. The value must be 0.
Reservation Identification	Any value in this field is ignored by the drive.
Extent List Length	Any value in this field is ignored by the drive.

NOTE: The drive returns RESERVATION CONFLICT status if the unit is already reserved for another initiator. Once the reservation is honored, any command from another initiator (except INQUIRY or REQUEST SENSE) is rejected with RESERVATION CONFLICT status.

5.36 RESERVE (10) Command (56h)

The RESERVE and the RELEASE commands are used for contention resolution in multiple-initiator systems. The RESERVE command is used to reserve a logical unit. If RESERVE (10) is used, then RELEASE (10) is also used.

Bit	7	6	5	4	3	2	1	0
Byte								
0	Operation Code (56h)							
1	Reserved			3rdPty	Reserved		LongID	Extent
2	Reservation Identification							
3	Third Party Device ID							
4 - 6	Reserved							
7 - 8	(MSB) Parameter List Length							(LSB)
9	Control							

Figure 5-90 RESERVE (10) Command Descriptor Block — Data Format

RESERVE (10) Command (56h) (continued)

Table 5-92 RESERVE (10) Command — Field Descriptions

Field	Description
3rd Pty	When set to 1, it indicates that the Third Party Device ID field is valid.
LongID	When 0, indicates that the third party device associated with the reservation release has a number smaller than 255 and the ID value can be sent within the CDB. If set = 1, indicates that the third party device ID is greater than 255, the ID value within the CDB is ignored, and the parameter list length is at least eight.
Third Party Device ID	<p>Required and used only when the 3rdPty bit is set, in which case this field specifies the SCSI ID of the initiator to be granted the reservation of the logical unit. The drive preserves the reservation until one of the following occurs:</p> <ul style="list-style-type: none"> • It is superseded by another valid RESERVE command from the initiator. • It is released by the same initiator. • It is released by a TARGET RESET message from any initiator. • It is released by a hard reset condition. <p>The drive ignores any attempt to release the reservation made by any other initiator. For example, if ID7 sends ID2 a Third Party reservation on behalf of ID6 (the target at ID2 gets reserved for the initiator ID6), then only ID7 can release the target at ID2 (using a Third Party release). ID6 cannot release the reservation even though the reservation was made on its behalf.</p>
Extent	<p>If 0, requests that the entire logical unit be reserved for use only by the initiator until the request is supplanted by another valid RESERVE command from the same initiator, or until released via a RELEASE (10) command from the reserving initiator, a hard reset, or a power on cycle.</p> <p>If set = 1, the extent reservation option is implemented. This option allows an application client in a multitasking environment to have multiple reservations.</p> <p>The size of the extent list is defined by the contents of the Parameter List Length field. The extent list consists of zero or more descriptors. Each descriptor defines an extent beginning at the specified logical block address for the specified number of blocks. If the number of blocks is 0, the extent begins at the specified logical block address and continues through the last logical block address on the logical unit. The data format of extent descriptors is shown in Figure 5-91.</p>

RESERVE (10) Command (56h) (continued)

Bit	7	6	5	4	3	2	1	0
Byte								
0	Reserved					RelAdr	Reservation Type	
1 - 3	(MSB)	Number of Blocks						(LSB)
4 - 7	(MSB)	Logical Block Address						(LSB)

Figure 5-91 Extent Descriptors — Data Format

The Reservation Type field defines the type of reservation in effect for the extent being reserved. Available types of reservations are shown in Table 5-93.

Table 5-93 Reservation Types

Reservation Type	Description
00b	READ shared. No write operations are permitted by any initiator to the reserved extent. Any initiator may read from the extent.
01b	WRITE exclusive. No other initiator may write to the indicated extent. Any initiator may read from the extent.
10b	READ exclusive. While this reservation is active, no other initiator may be permitted to read to the indicated extent.
11b	Exclusive access. While this reservation is active, no other initiator is permitted any access to the reserved extent.

RelAdr is always 0: The Quantum Atlas 10K II hard disk drive does not support relative addressing.

If both the LongID and the Extent bits = 1, then the parameter list takes the format shown in Figure 5-92.

Bit	7	6	5	4	3	2	1	0
Byte								
0 - 7	(MSB)	Third Party Device ID						(LSB)
8 - n	(MSB)	Extent Descriptors (See Figure 5-91 Above)						(LSB)

Figure 5-92 Parameter List When LongID and Extent Bits = 1 — Data Format

RESERVE (10) Command (56h) (continued)

If the LongID bit = 1 and the Extent bit = 0, then the parameter list length is eight and the parameter list has the format shown in Figure 5-93.

Bit	7	6	5	4	3	2	1	0
Byte								
0 - 7	Third Party Device ID							(LSB)

Figure 5-93 RESERVE (10) ID Only Parameter List — Data Format

5.37 REZERO UNIT Command (01h)

The REZERO UNIT command sends the heads to cylinder 0 (Logical Block Address 0). Once the heads are on the track, they are not guaranteed to stay at cylinder 0; periodic calibrations may send them to a different cylinder. If periodic calibrations have been disabled by using the MAN bit (byte 3, bit 6) of the Vendor Unique Function Control Page in the MODE SELECT command (15h), the REZERO UNIT command will cause the drive to perform the next scheduled calibration.

Bit	7	6	5	4	3	2	1	0
Byte								
0	Operation Code (01h)							
1 - 4	Reserved							
5	Control							

Figure 5-94 REZERO UNIT Command Descriptor Block — Data Format

5.38 SEEK (6) Command (0Bh)

The SEEK (6) command uses a 6-byte Command Descriptor Block to request the drive to send the drive heads to the requested Logical Block Address. The drive responds by placing the heads at the correct Logical Block Address. When the heads are at the specified Logical Block Address, they are not guaranteed to stay in that position since periodic calibrations may send them to a different cylinder.

Bit	7	6	5	4	3	2	1	0
Byte								
0	Operation Code (0Bh)							
1	Reserved			Logical Block Address				
2 - 3	Logical Block Address							
4	Reserved							
5	Control							

Figure 5-95 SEEK (6) Command Descriptor Block — Data Format

5.39 SEEK (10) Command (2Bh)

The SEEK (10) command uses a 10-byte Command Descriptor Block to request the drive to send the drive heads to the requested Logical Block Address. The drive responds by placing the heads at the correct Logical Block Address. When the heads are at the specified Logical Block Address, they are not guaranteed to stay in that position since periodic calibrations may send them to a different cylinder.

Bit	7	6	5	4	3	2	1	0
Byte								
0	Operation Code (2Bh)							
1	Reserved							
2 - 5	Logical Block Address							
6 - 8	Reserved							
9	Control							

Figure 5-96 SEEK (10) Command Descriptor Block — Data Format

5.40 SEND DIAGNOSTIC Command (1Dh)

The SEND DIAGNOSTIC command sends a diagnostic page to the drive. Refer to the RECEIVE DIAGNOSTICS RESULTS command for a description of the diagnostic pages and the use of the diagnostic commands. The Supported Diagnostic Page List (Figure 5-98) and Translate Page List (Figure 5-99) can be sent with this command

Bit	7	6	5	4	3	2	1	0
Byte								
0	Operation Code (1Dh)							
1	Reserved			PF	Rsvd	Self Test	DevOnl	UnitOfl
2	Reserved							
3 - 4	Parameter List Length							
5	Control							

Figure 5-97 SEND DIAGNOSTIC Command Descriptor Block — Data Format

Table 5-94 SEND DIAGNOSTIC Command — Field Descriptions

Field	Description
PF	Page Format. When the PF bit is 1, the parameters that follow conform to the Supported Diagnostic Page List (Figure 5-96) or the Translate Address Page (Figure 5-97). SelfTest is ignored when PF = 1. The PF bit must be 1 if Parameter List Length is not 0.
SelfTest	If the PF bit is 0 and the SelfTest bit is 1, the drive runs its self-test routine. GOOD status is returned to the initiator if the self-test passes. CHECK CONDITION status is returned along with any available sense data if the self-test fails. This field is ignored if the PF bit is 1.
DevOfl	Any value in this field is ignored by the drive.
UnitOfl	Any value in this field is ignored by the drive.
Parameter List Length	When the Parameter List Length field is non-zero, the value specifies the length (in bytes) of the parameter list that is to be transferred from the initiator to the drive during the command's data-out buffer transfer. This field is only valid when the PF bit = 1 and must be 0 when the PF bit is 0.

*SEND DIAGNOSTIC Command (1Dh) (continued)***5.40.1 Supported Diagnostic Page List**

Figure 5-98 shows the Supported Diagnostic Page List that can be supplied with a SEND DIAGNOSTIC command to request that the Supported Diagnostic Page List be returned after the next RECEIVE DIAGNOSTIC RESULTS command.

Bit Byte	7	6	5	4	3	2	1	0
0	Page Code (00h)							
1	Reserved							
2 - 3	Page Length (00h)							

Figure 5-98 Supported Diagnostic Page List — Data Format

5.40.2 Translate Address Page

Figure 5-99 shows the Translate Address Page that can be supplied with a SEND DIAGNOSTIC command to request that the Translate Address Page be returned after the next RECEIVE DIAGNOSTIC RESULTS command.

Bit Byte	7	6	5	4	3	2	1	0
0	Page Code (40h)							
1	Reserved							
2 - 3	Page Length (0Ah)							
4	Reserved					Supplied Format 000b (LBA)		
5	Reserved					Translate Format 101b (Physical Sector)		
6 - 9	Address to Translate (LBA Format)							
10 - 13	Reserved							

Figure 5-99 Translate Address Page — Data Format

5.41 SET DEVICE IDENTIFIER Command (A4h)

The SET DEVICE IDENTIFIER command requests that the device identifier information in the logical unit be set to the value send via the SET DEVICE IDENTIFIER command's parameter list.

Upon successful completion of a SET DEVICE IDENTIFIER command, a Unit Attention is generated for all initiators except the one that issued the command.

Bit	7	6	5	4	3	2	1	0
Byte								
0	Operation Code (A4h)							
1	Reserved			Service Action (06h)				
2 – 5	Reserved							
6 – 9	(MSB) Parameter List Length (LSB)							
10	Reserved							
11	Control							

Figure 5-100 SET DEVICE IDENTIFIER Command Descriptor Block — Data Format

Table 5-95 SET DEVICE IDENTIFIER Command Descriptor Block — Field Descriptions

Field	Description
Service Action	Must = 06h. Any other value forces Check Condition, Illegal Request.
Parameter List Length	This field specifies the length, in bytes, of the Identifier to be transferred from the application client to the device server. The maximum value for this field is 64 bytes. A parameter list length of 0 indicates that no data will be transferred, and that subsequent REPORT DEVICE IDENTIFIER commands will return an Identifier length of 0. If the parameter list length exceeds 64 bytes, then the drive returns a Check Condition status with the sense key set to Illegal Request, and an additional sense code of Invalid Field in CDB.

SET DEVICE IDENTIFIER Command (A4h) (continued)

The SET DEVICE IDENTIFIER command's parameter list contains Identifier to be set by the logical unit to which the command is sent.

Bit	7	6	5	4	3	2	1	0
Byte								
0 - n	(MSB)	Identifier						(LSB)

Figure 5-101 SET DEVICE IDENTIFIER Parameter List — Data Format

Table 5-96 SET DEVICE IDENTIFIER Parameter List —Field Descriptions

Field	Description
Identifier	The value reported in this field is the data to be returned by all subsequent REPORT DEVICE IDENTIFIER commands, until a new SET DEVICE IDENTIFIER command is issued that changes the data.

5.42 START STOP UNIT Command (1Bh)

The START STOP UNIT command enables or disables the logical unit for media access operations. This command gives the initiator remote control over drive spin-up and spin-down.

A jumper (or backplane connection) is used to select the spin-up mode. When configured for remote control of spin-up (Option 3), the drive requires a START STOP UNIT Command to spin up. See Chapter 3 for jumper-programming instructions for your specific hard disk drive.

Bit	7	6	5	4	3	2	1	0
Byte								
0	Operation Code (1Bh)							
1	Reserved							Immed
2 - 3	Reserved							
4	Power Conditions				Reserved		LoEj	Start
5	Control							

Figure 5-102 START STOP UNIT Command Descriptor Block — Data Format

Table 5-97 START STOP UNIT Command — Field Descriptions

Field	Description
Immed	Immediate. When 1, the drive returns GOOD status immediately after validating the Command Descriptor Block, but before the drive spins up. When 0, the drive does not return GOOD status until after the drive has spun up successfully and is ready for media access commands.
LoEj	Load Eject. This value must be zero, because the disk drive contains non-removable medium.
Start	When Start is 1, the unit spins up (if it is not currently spinning) and is left in a state where subsequent media access commands can succeed. When the value is 0, the drive spins down. Subsequent media access commands are rejected with CHECK CONDITION status and the proper sense data.
Power Conditions	The Power Conditions field is not supported.

5.43 SYNCHRONIZE CACHE Command (35h)

The SYNCHRONIZE CACHE command ensures that all logical blocks, within the specified range in the drive's cache memory, have their most recent data value recorded on the physical medium. If more recent data for a logical block, within the specified range, exists in the cache memory than on the physical medium, then the logical blocks from the cache memory is written to the physical medium. Logical blocks are not necessarily removed from the cache as a result of the Synchronize Cache operation.

Bit	7	6	5	4	3	2	1	0
Byte								
0	Operation Code (35h)							
1	Reserved						Immed = 0	RelAdr
2 - 5	Logical Block Address							
6	Reserved							
7 - 8	Number of Blocks							
9	Control							

Figure 5-103 SYNCHRONIZE CACHE Command Descriptor Block — Data Format

Table 5-98 SYNCHRONIZE CACHE Command — Field Descriptions

Field	Description
Immed	Immediate. Must be 0. The drive returns status after the synchronize cache operation has been completed.
RelAdr	Relative Address. Must be 0. Relative addressing is not supported.
Logical Block Address	The address of the first logical sector of the range to be operated upon by this command.
Number of Blocks	Specifies the total number of contiguous logical blocks within the range. When set to 0, the drive removes all sectors from the given Logical Block Address on. When Number of Blocks and Logical Block Address are both 0, the drive removes all unwritten cache sectors to the medium. It is not considered an error when no sectors within the specified range are in the cache memory.

5.44 TEST UNIT READY Command (00h)

The TEST UNIT READY command returns GOOD status if the drive is ready to receive medium access commands (that is, READ, WRITE, VERIFY, etc.). If the drive is not ready to receive medium access commands, it returns CHECK CONDITION status with the appropriate sense keys and qualifiers.

Bit	7	6	5	4	3	2	1	0
Byte								
0	Operation Code (00h)							
1 - 4	Reserved							
5	Control							

Figure 5-104 TEST UNIT READY Command Descriptor Block — Data Format

5.45 VERIFY Command (2Fh)

The VERIFY command verifies the data on the medium for the Logical Block Address range specified.

Bit	7	6	5	4	3	2	1	0
Byte								
0	Operation Code (2Fh)							
1	Reserved			DPO	Reserved		BytChk	RelAdr
2 - 5	Logical Block Address							
6	Reserved							
7 - 8	Verification Length							
9	Control							

Figure 5-105 VERIFY Command Descriptor Block — Data Format

Table 5-99 VERIFY Command — Field Descriptions

Field	Description
DPO	Disable Page Out. When the value of the field = 1, the device server shall assign the logical blocks accessed by this command the lowest priority for being fetched into or retained by the cache. When DPO = 0, the priority shall be determined by fields in the Caching page.
BytChk	Byte Check. When the value is 1, the drive performs a byte-by-byte comparison of the data on the medium with the data sent by the initiator in the data-out buffer transfer. When 0, the drive performs a medium verification with no data comparison. In this case, no data is sent by the initiator.
RelAdr	Relative Address. This feature is not supported; the field must be 0.
Logical Block Address	The address of the first logical sector to be verified.
Verification Length	The number of contiguous sectors that are to be verified. The drive returns a CHECK CONDITION status with the correct sense data if the specified transfer extends beyond out of the logical sector space.

5.46 WRITE (6) Command (0Ah)

The WRITE (6) command requests that the drive write initiator-supplied data to the medium.

The Command Descriptor Block for the WRITE (6) command is shown in the figure below; the table that follows describes the data fields.

Bit	7	6	5	4	3	2	1	0
Byte								
0	Operation Code (0Ah)							
1	Reserved			Logical Block Address				
2 - 3	Logical Block Address							
4	Transfer Length							
5	Control							

Figure 5-106 WRITE (6) Command Descriptor Block — Data Format

Table 5-100 WRITE (6) Command — Field Descriptions

Field	Description
Logical Block Address	The address of the first logical block where the Write operation will occur.
Transfer Length	The number of contiguous logical blocks of data to transfer. A transfer length of 0 indicates 256 logical blocks are to be transferred. Any other value indicates the actual number of logical blocks that will be transferred.

5.47 WRITE (10) Command (2Ah)

The WRITE (10) command requests that the drive write initiator-supplied data to the medium.

Bit	7	6	5	4	3	2	1	0
Byte								
0	Operation Code (2Ah)							
1	Reserved			DPO	FUA	Reserved		RelAdr
2 - 5	Logical Block Address							
6	Reserved							
7 - 8	Transfer Length							
9	Control							

Figure 5-107 WRITE (10) Command Descriptor Block — Data Format

Table 5-101 WRITE (10) Command — Field Descriptions

Field	Description
DPO	<p>Disable Page Out. When the value of the field is 1, the drive tells the controller to bypass cache as much as possible. The buffers used for the WRITE data are given a low priority. This bit may be useful if the host CPU is performing a background copy operation and does not want to wipe out the drive's cache.</p> <p>When DPO is 0, the drive tells the controller to cache this WRITE data in a normal manner.</p> <p>NOTE: The drive normally assigns a low priority to write caching. Therefore, the DPO bit in a WRITE command is ignored.</p>
FUA	<p>Force Unit Access. When 1, the data must be written on the medium before GOOD status is returned. When 0, GOOD status may be returned after the data is fetched from the initiator and before it is written to the medium. If the WCE (Write Cache Enable) bit on the MODE SELECT command Caching Page is clear, the FUA bit is ignored and behaves as if it were always set (that is, if WCE = 0, the controller forces FUA = 1).</p>
Logical Block Address	The address of the first logical block where the Write operation will occur.
RelAdr	Relative Address. This drive does not support this feature; the field must be 0.
Transfer Length	The number of contiguous logical blocks of data to write. The drive returns CHECK CONDITION status with the correct sense data if the specified transfer would cause the write operation to extend beyond the Logical Block Address space.

5.48 WRITE AND VERIFY Command (2Eh)

The WRITE AND VERIFY command writes host-supplied data to the unit and verifies it was written correctly.

Bit	7	6	5	4	3	2	1	0
Byte								
0	Operation Code (2Eh)							
1	Reserved			DPO	Reserved		BytChk	RelAdr
2 - 5	Logical Block Address							
6	Reserved							
7 - 8	Transfer Length							
9	Control							

Figure 5-108 WRITE AND VERIFY Command Descriptor Block — Data Format

Table 5-102 WRITE AND VERIFY Command — Field Descriptions

Field	Description
DPO	<p>Disable Page Out. When the value of the field is 1, the drive tells the controller to bypass cache as much as possible. The buffers used for the WRITE data are given a low priority. This bit may be useful if the host CPU is performing a background copy operation and does not want to wipe out the drive's cache.</p> <p>When DPO is 0, the drive tells the controller to cache this WRITE data in a normal manner.</p> <p>NOTE: The drive normally assigns a low priority to write caching. Therefore, the DPO bit in a WRITE command is essentially ignored.</p>
BytChk	<p>Byte Check. When 1, the drive performs a byte-by-byte comparison of the data written on the media with the data transferred from the initiator.</p> <p>When 0, the drive performs a medium verification with no data comparison of the sectors written by the command.</p>
RelAdr	Relative Address. This feature is not supported and the field must be 0.

(continued)

WRITE AND VERIFY Command (2Eh) (command)

Table 5-102 WRITE AND VERIFY Command — Field Descriptions (continued)

Field	Description
Logical Block Address	The address of the first logical block where the WRITE operation will occur.
Transfer Length	The number of contiguous logical blocks of data to write and verify. The drive returns CHECK CONDITION status with the correct sense data if the specified transfer would cause the write operation to extend beyond the Logical Block Address space.

5.49 WRITE BUFFER Command (3Bh)

The WRITE BUFFER and the READ BUFFER commands allow the initiator to gain access to the data storage areas within the drive controller. These commands are used to check the integrity of the buffer-initiator data path independent of the storage medium, typically for fault isolation. The commands can be used to verify the caching policies of the controller or other kinds of debug work. In addition, the WRITE BUFFER command supports controller microcode downloads.

Both the READ BUFFER and WRITE BUFFER commands have a *Buffer ID* and a *Buffer Offset* data field defined in their Command Descriptor Block. In combined Header and Data Mode, if both are set to 0, reference is made to a single, 512-byte buffer set aside in the controller's data memory, independent of cache. This allows an initiator to verify the data path from the controller's data memory to the initiator without affecting the drive's internal data caches.

Bit	7	6	5	4	3	2	1	0
Byte								
0	Operation Code (3Bh)							
1	Reserved				Mode			
2	Buffer ID							
3 - 5	Buffer Offset							
6 - 8	Parameter List Length							
9	Control							

Figure 5-109 WRITE BUFFER Command Descriptor Block — Data Format

WRITE BUFFER Command (3Bh) (continued)

Table 5-103 WRITE BUFFER Command — Field Descriptions

Field	Description
Mode	Four modes are supported by the drive: Combined Header and Data, Data Only, Download Microcode and Save, and Write Data from Echo Buffer. The permissible values for Buffer ID, Buffer Offset are as follows:
0000b	<p>Combined Header and Data. – In this mode, the data to be transferred is preceded by a 4-byte header that contains reserved bytes that must be set to 0. The resulting descriptor and data are fetched from the initiator during the command's data-out buffer transfer and, after the descriptor has been validated, only the data is written to the controller's memory.</p> <p>Buffer ID must be 0</p> <p>Buffer Offset must be 0</p> <p>Parameter List Length must be set to between 0 and 516 (4 bytes of header and up to 512 bytes of data).</p>
0010b	<p>Data Only.– In this mode, the data-out buffer transfer contains buffer data. Data is written to the drive's buffer starting at the location specified by the Buffer ID and Buffer Offset.</p> <p>Buffer ID. When non-zero, the value specifies a specific track line in the controller's cache. If Buffer ID of zero is specified, the buffer referenced is the same as if Buffer ID of one were specified.</p> <p>Buffer Offset. When non-zero, the value can range from 0 to one less than the number of sectors/track multiplied by 512. This value must be either a zero or a multiple of 512.</p> <p>Parameter List Length. This value must be a multiple of 512. The maximum value for this field is one more than the number of sectors/track times 512. The minimum value is 0.</p>
0101b	<p>Download Microcode and Save In this mode, a 256 KB microcode image is transferred to the drive and, if valid, is saved in the controller's non-volatile memory. This image can be transferred all at once in one 256K buffer, or it can be transferred in 32 separate 8K buffers. The command's status phase indicates the success or failure of the code load operation: GOOD status means the new microcode has been loaded and saved; CHECK CONDITION status means that the original microcode is still in force. CHECK CONDITION status is returned if the supplied microcode image is invalid. Sense data reporting the reason for the failure is not preserved.</p> <p>The controller does a hard reset on command completion. In Download Microcode and Save mode, the other Command Descriptor Blocks are supported as follows:</p> <p>Buffer ID – Undefined and ignored.</p> <p>Buffer Offset – Undefined and ignored for 256 KB transfer. For a transfer of 32 separate 8K buffers, Buffer Offset represents the offset of the buffer being transferred (0, 8192, 16384, ...).</p> <p>Parameter List Length – Must be 8192 or 262144.</p>
1010b	<p>Write Data from Echo Buffer</p> <p>In this mode, the host transfers data and stores it in an echo buffer. The data will be preserved in the echo buffer unless there is an intervening command.</p> <p>Buffer ID Ignored.</p> <p>Buffer Offset Ignored.</p>

5.50 WRITE LONG Command (3Fh)

The WRITE LONG and READ LONG commands allow the initiator to issue raw reads and writes of physical sectors addressed as logical blocks with an assumed sector length of 512 bytes. These commands are used to verify the correction capability of the drive's ECC and to create defects of a known kind and location to verify the drive's recovery and bad sector handling policies.

Bit	7	6	5	4	3	2	1	0
Byte								
0	Operation Code (3Fh)							
1	Reserved							RelAdr
2 - 5	Logical Block Address							
6	Reserved							
7 - 8	Byte Transfer Length (562)							
9	Control							

Figure 5-110 WRITE LONG Command Descriptor Block — Data Format

Table 5-104 WRITE LONG Command — Field Descriptions

Field	Description
RelAdr	Relative Address. This drive does not support this feature; the field must be 0.
Logical Block Address	The address of the first logical sector where the Write operation will occur.
Byte Transfer Length	<p>Must be 562. This value reflects the number of bytes spanned by the ECC, including the ECC itself, rounded up to the next word (this is easiest understood in bits: 4096 data bits + 1 force error bit + 15 Logical Block Address bits + 16 cross check bits + 2 pad bits + 360 ECC bits + 6 pad bits that are not written to media but are only needed for transfer = 4496 bits or 562 bytes). Any other value in this field results in a CHECK CONDITION status and the appropriate sense data.</p> <p>The initiator should first issue a READ LONG, then modify the data portion only, and then issue a WRITE LONG. It is also recommended that the drive be reformatted after these commands are used.</p> <p>If the drive is configured with a sector size other than 512 bytes per sector, the WRITE LONG command affects the first 512 byte sector set aside by the drive for the given Logical Block Address.</p>

5.51 WRITE SAME Command (41h)

The WRITE SAME command requests that the drive write the single sector of data transferred by the initiator to the medium multiple times.

Bit	7	6	5	4	3	2	1	0
Byte								
0	Operation Code (41h)							
1	Reserved					PBdata	LBdta	RelAdr
2 - 5	Logical Block Address							
6	Reserved							
7 - 8	Number of Blocks							
9	Control							

Figure 5-111 WRITE SAME Command Descriptor Block — Data Format

Table 5-105 WRITE SAME Command — Field Descriptions

Field	Description
PBdata	Physical Block Data. – Not supported. This field must be zero.
LBdata	Logical Block Data. – A value of 1 requests that the drive replace the first four bytes of the data to be written to the current logical sector with the logical sector address of the sector currently being written.
RelAdr	Relative Address. – Not supported. This field must be zero.
Logical Block Address	The address of the first logical sector where the Write operation will occur.
Number of Blocks	This field contains the number of contiguous logical blocks that are to be written. A value of 0 requests that all the remaining logical blocks on the medium are to be written.

5.52 WRITE SKIP MASK Command (EAh)

The WRITE SKIP MASK command allows the initiator to condition the function of the WRITE commands to allow only selected data to be transferred over the SCSI interface. A WRITE SKIP MASK command precedes a WRITE (10) command; the WRITE (10) command must be linked to the WRITE SKIP MASK command or the skip mask is lost and a CHECK CONDITION status is returned.

Bit	7	6	5	4	3	2	1	0
Byte								
0	Operation Code (EAh)							
1	Reserved							
2 - 5	(MSB)	Logical Block Address						(LSB)
6	Skip Mask Length							
7 - 8	(MSB)	Transfer Length						(LSB)
9	VU	Reserved					Flag	Link

Figure 5-112 WRITE SKIP MASK Command Descriptor Block — Data Format

Table 5-106 WRITE SKIP MASK Command — Field Descriptions

Field	Description
Logical Block Address	This field specifies the first logical block of the corresponding WRITE (10) COMMAND. This allows a SEEK to begin immediately. The LBA of the following WRITE (10) command must be the same as the LBA of the WRITE SKIP MASK command.
Skip Mask Length	This field indicates the length in bytes of the skip mask. The mask is up to 2048 bits, giving a data address capability of 1 MB in 512-byte block length. The mask consists of a sequence of bits where a "1" bit is for the blocks of data to be transferred, and a "0" for the blocks of data that are to be skipped. The first "1" bit of the mask corresponds to the first LBA to be transferred. Any unused bits at the end of the last byte in the mask must be set to 0. The mask is transferred in the data-out buffer transfer of the skip mask command.
Transfer Length	This field must be equal to the transfer length in the linked WRITE (10) command. The transfer length field must be equal to the number of "1" bits in the skip mask. If a WRITE (10) is used with 0 as the transfer length, a transfer length of 256 data blocks is implied.

5.53 XDREAD (10) Command (52h)

The XDREAD (10) command requests that the target transfer to the initiator the XOR data generated by an XDWRITE (10) command. The XOR data transferred is identified by Logical Block Address and transfer length that are the same as those specified in a prior XDWRITE (10) command. If a match is not found, the command is terminated with a CHECK CONDITION status. The sense data is set to ILLEGAL REQUEST: INVALID FIELD IN CDB.

Bit	7	6	5	4	3	2	1	0
Byte								
0	Operation Code (52h)							
1	Reserved							
2 - 5	Logical Block Address							
6	Reserved							
7 - 8	Transfer Length							
9	Control							

Figure 5-113 XDREAD (10) Command Descriptor Block — Data Format

Table 5-107 XDREAD (10) Command — Field Descriptions

Field	Description
Logical Block Address	Specifies the address of the prior XDWRITE (10) command that generated the XOR data in the buffer that is being requested.
Transfer Length	Specifies the number of contiguous logical blocks of data to transfer. This length must match the length of the original XDWRITE (10) command that generated the XOR data requested.

5.54 XDWRITE (10) COMMAND (50h)

The XDWRITE (10) command requests that the target XOR the data transferred with the data on the medium. The resulting XOR data is stored in the target's buffer. The disposition of the data transferred from the initiator is controlled by the Disable Write bit.

The resulting XOR data remains in the target's buffer until it is retrieved by an XDREAD command with starting Logical Block Address and transfer length fields that match the starting Logical Block Address and transfer length of the XDWRITE command.

Bit	7	6	5	4	3	2	1	0
Byte								
0	Operation Code (50h)							
1	Reserved			DPO	FUA	Disable Write	Reserved	
2 - 5	Logical Block Address							
6	Reserved							
7 - 8	Transfer Length							
9	Control							

Figure 5-114 XDWRITE (10) Command Descriptor Block — Data Format

Table 5-108 XDWRITE (10) Command — Field Descriptions

Field	Description
DPO	Disable Page Out. When the value = 0, the drive caches the data read, using its normal Logical Read Unit Policies.
	When the value = 1, the drive assigns to logical blocks accessed the lowest priority for being fetched into or retained by the cache.
FUA	Force Unit Access. When the value is 0, the drive is allowed to satisfy the read with cache-resident data.
	A value of 1 indicates that the drive accesses the media in performing the command.
Disable Write	A Disable Write bit of 0 indicates that the data transferred from the initiator is to be written to the medium after the XOR operation is complete. A Disable Write bit of 1 indicates that the data will not be written to the medium. If the Disable Write bit is set to 1, the FUA bit is ignored.

(continued)

XDWRITE (10) COMMAND (50h) (continued)

Table 5-108 XDWRITE (10) Command — Field Descriptions (continued)

Field	Description
Logical Block Address	Specifies the address of the first logical block where the WRITE operation will occur.
Transfer Length	Specifies the number of contiguous logical blocks to write. The drive returns CHECK CONDITION status with the correct sense data if the specified transfer will cause the WRITE operation to extend beyond the Logical Block Address space.

5.55 XPWRITE (10) COMMAND (51h)

The XPWRITE (10) command requests that the target XOR the data transferred with the data on the medium then writes the XOR data to the medium.

The Logical Block Address field specifies the Starting Logical Block address for the target to read data from its medium. It also specifies the starting Logical Block Address at which to write the XOR result to its medium.

Bit	7	6	5	4	3	2	1	0
Byte								
0	Operation Code (51h)							
1	Reserved			DPO	FUA	Reserved		
2 - 5	Logical Block Address							
6	Reserved							
7 - 8	Transfer Length							
9	Control							

Figure 5-115 XPWRITE (10) Command Descriptor Block — Data Format

Table 5-109 XPWRITE (10) Command — Field Descriptions

Field	Description
DPO	Disable Page Out. When the value = 0, the drive caches the data read, using its normal Logical Read Unit Policies. When the value = 1, the drive assigns to logical blocks accessed the lowest priority for being fetched into or retained by the cache.
FUA	Force Unit Access. When the value is 0, the drive is allowed to satisfy the read with cache-resident data. A value of 1 indicates that the drive accesses the media in performing the command.

(continued)

XPWRITE (10) Command (51h) (continued)

Table 5-109 XPWRITE (10) Command — Field Descriptions (continued)

Field	Description
Logical Block Address	Specifies the address of the first logical block from which to read data from its medium and the first logical sector to which to write the XOR data back to its medium.
Transfer Length	Specifies the number of contiguous logical blocks to write. The drive returns CHECK CONDITION status with the correct sense data if the specified transfer will cause the WRITE operation to extend beyond the Logical Block Address space.

This chapter provides descriptions of many of the features incorporated in the Quantum Atlas 10K II Ultra 160/m SCSI hard disk drive family.

6.1 KEY FEATURES

Atlas 10K II Disk Drives provide unequalled quality and performance. They are designed to improve the performance of high-end workstations, file servers, and storage systems. Among their features are the following:

- *9.2, 18.4, 36.7, or 73.4 GB formatted capacity (512 bytes/block)*
- *MR heads and banded recording for higher overall bit-densities*
- *Hardware XOR to automatically detect and correct errors and defects in the data stream for array applications (RAID). Third Party XOR support for Force.*
- *Self-Monitoring, Analysis, and Reporting Technology (S.M.A.R.T.) Revision 2 support*
- *ORCA (Optimized Reordering Command Algorithm) and tagged command queuing for greater throughput*
- *SCAM-2 (SCSI Configured AutoMatically) Auto ID configuration*
- *8 MB Track-oriented segmented cache buffer with prefetch*
- *Automatic Power Management System with power-saving sub-states*
- *10,000 rpm rotational speed yielding 3.0 ms average latency*
- *4.7 ms average random seek time (READ) and a TBD ms average access-to-data time for 1.0-inch form factor drives; 5.2 ms second average random seek time (READ) and TBD ms average access-to-data time for the 1.6-inch form factor drive.*
- *Embedded servo system for exceptional head positioning accuracy and long life*
- *“Zero latency” reads and writes reduce data access time*
- *Full SCSI-3 compliance ¹ (compatible with SCSI-2 and SCSI-1)*

¹ The drives are in compliance with the SCSI-3 draft proposed ANSI standards as they exist at the time of product release; SCSI-3 features may not be enabled on drives when shipped.

- *Advanced Ultra3 160/m-LVD, Ultra2-LVD, and Ultra-SE with 68-pin wide connector or 80-pin SCA-2 connector*
- *Superior data integrity:*
 - *352-bit Reed-Solomon quad-burst ECC provides fast on-the-fly correction for small errors and robust software-based correction for larger errors*
 - *EDC internal protection from SCSI bus to media*

Many of these features are described in this chapter.

The drive contains a large, segmented buffer in RAM that is used to maximize the throughput for sequential read streams. This cache is organized as a number of separate track line segments. Each track line segment can hold data from one complete track on the media, temporarily assuming the identity of the track for all input/output functions.

The buffer is also used to combine or coalesce I/O data on a track-by-track basis, prior to accessing the disk media. Performance is greatly improved by gathering and combining the data on a track basis rather than on a command basis. With this unique system, multiple commands for a particular track are coalesced in the buffer before accessing the media.

As a result, multiple I/O commands, that address a single track are logically combined to avoid multiple disk accesses on multiple rotations. Since most reads and writes are sequential, this greatly increases throughput for most applications.

6.2 PREFETCH

The drive supports prefetch. After the controller has copied an entire media track into a cache track line segment, a unique algorithm determines whether:

- The most recently used cache track line (track n) satisfies a read request;
- The cache memory also includes another track line (track n-1) corresponding to the media track preceding the one that was just copied, and;
- There is no other work currently waiting for service by the read/write system.

If all three of these conditions are met, there is a high probability that the host will call for data from the next media track (track n+1). So, the drive controller automatically issues a read for that data and substitutes it for the information in the n-1 track line segment in cache. This process, called prefetch, gives the drive a very high cache hit rate for sequential reads.

6.3 WRITE-BACK CACHING

The drive supports write-back caching. In this case, writes are considered complete when the data is loaded into the cache, before it is written to the media. This allows the host to go on to other tasks, while the drive assumes responsibility for the data and ensuring that it is written to media as soon as possible. If write-back caching is employed, it is recommended that an uninterruptible power source (UPS) be provided for optimal data security.

6.4 ZERO LATENCY READ/WRITE

An average of half a revolution of latency is saved by starting to read or write as soon as a seek settles on the designated track. If reading, the initial data read is stored in a buffer (cache) until the actual starting address data block is reached and read. The starting address data is then transferred to the bus as the rest of the data is transferring from media to the buffer. When the initial data address read is reached again, that data in the buffer is transferred to the bus. This ensures that the track data goes out in media order and eliminates the rotational latency associated with waiting for the starting address to rotate to the head.

Writing works essentially the same way only in reverse. The data is first transferred to the drive's buffer. After seeking to the desired track, a write pointer is set to the current location of the head on the track and writing is commenced from the buffer, starting at the corresponding place in the data.

6.5 DISCONNECT-RECONNECT

System throughput can be improved by disconnecting the drive from the initiator during physical positioning operations, thereby freeing up the SCSI bus for other I/O processes. After the drive has determined that there will be a delay, it disconnects itself from the SCSI bus by sending a DISCONNECT message and enters the BUS FREE phase.

When the drive is ready to resume data transfer, it arbitrates for the SCSI bus and, after winning arbitration, reconnects to the initiator and sends an IDENTIFY message via the MESSAGE IN phase. This revives the I_T_L nexus (initiator-target-logical unit connection) so the initiator can retrieve the correct set of data pointers for the I/O process. The initiator restores the active pointers to their most recent saved values, and the drive continues to finish the original I/O process.

Disconnect-Reconnect is controlled with the Mode Select Command (15h) for the Disconnect/Reconnect (02h) mode page, as described in Chapter 6.

6.6 TRACK AND CYLINDER SKEWING

The disk drive improves data throughput by skewing track and cylinder addresses. When the drive switches heads or tracks, or both, to access sequential data, the rotation of the disk media allows one or more physical blocks to pass by the read/write head before the head is ready to continue the transfer. Skewing adjusts the block addresses so that the head switch and settle times and the media rotation coincide to bring the head exactly to the next logical block. This minimizes rotational latency (and increases throughput) when data is accessed sequentially.

6.7 AVERAGE ACCESS TIME

A 10,000 rpm rotation speed yields an average latency of 3.00 ms.

6.8 EMBEDDED SERVO SYSTEM

Embedded servo information is written in a spoke configuration on every track, on every disk surface. The spokes (or headers) consist of quadrature analog patterns and digital address data. The digital portions of the spoke data are read and used to locate the desired track, spoke, and head number. The quadrature analog signal portion is detected and used by a servo feedback control loop to precisely position the head on the track center.

6.9 DATA INTEGRITY AND SECURITY

The disk drives use a combination of parity checking, error detection coding (EDC), error correction coding (ECC), and checkpointing to protect stored data from media errors, transfer or addressing errors, or errors introduced during block reallocation.

6.9.1 Media Error Protection

To ensure that data read is the same as data written, the drive computes and appends an Error Correction Code (ECC) to each block of data stored. The drive uses a 352-bit Reed Solomon code with a 4:1 interleave, which can correct up to 20 bytes in each block.

The drive can also correct up to 2 bytes per interleave (up to 8 per block) in hardware (“on-the-fly”), with no loss in throughput.

6.9.2 Transfer Error Protection

An end-to-end error detection code (EDC) protects data from any errors introduced by internal buses, the disk controller chip, the data cache, or the SCSI interface. An EDC is calculated and added to each data block as the data arrives from the SCSI bus (after SCSI bus parity is checked). The EDC is stored with the data and protected by the block ECC for added security. On reading or writing, the EDC is checked as the data is transferred between buffer RAM and the media or the SCSI bus.

6.9.3 Addressing Error Protection

Each data block on the media is identified and located by a servo spoke address. The spoke address consists of a two-byte word. Each spoke has multiple copies of the least significant bytes of the address. The disk hardware requires that a majority of the copies agree and that the result agrees with the expected head, track, and spoke number, before it will read or write the data.

To further protect against addressing errors, the logical address (LBA) of the data is added to the EDC of each block. If data is written to the wrong block and subsequently read, or read from the wrong block, the error will be flagged.

The hardware does not allow a blind read of a data block; the firmware must request specific data blocks. Even if the head selection hardware malfunctions, it is not possible for the drive to return data from the wrong head.

6.9.4 Data Sector Reallocation Error Protection

In any SCSI disk drive, bad blocks may be reallocated. However, a power failure or unrecoverable data could threaten data integrity during a block reallocation.

The reallocation and defect list storage algorithms prevent a reallocation from being lost due to a power failure. Once a reallocation starts, the information about the data block to be moved is stored on the media. As the reallocation progresses, checkpoint information about the reallocation (such as the new destination and the data in transit) is recorded. If the reallocation is interrupted, the process can automatically continue from the last checkpoint without data loss.

If a block reallocation is performed on *unrecoverable data* (rare), a Bad Data Mark is set in the new location. Thus, there is no chance of undetected bad data being generated from a reallocation. Bad block reallocation is discussed in section 6.16.3.

6.9.5 Data Verification

Data can be verified in two ways. First, the integrity of data in selected areas of the media can be checked using the VERIFY command. Secondly, the WRITE AND VERIFY command writes host-supplied data and then verifies that it was written correctly.

The Verify Error Recovery Page of the MODE SELECT COMMAND controls the drive's response to error conditions that arise during the VERIFY command and during the verify operation of the WRITE AND VERIFY command.

6.10 TAGGED COMMAND QUEUING

The drive supports all three SCSI-defined queue tag messages. When queuing is enabled (default condition), the drive can accept any of the following:

- **Simple** - specifies that the command is to be placed in the drive's command task set (queue). If several commands are present, the drive may reorder them to increase overall throughput.
- **Head of Queue** - directs the drive to place the command at the beginning of the queue, to be executed next. Consecutive commands with Head of Queue specified are executed in last-in-first-out (LIFO) order. Queue depth is 64.
- **Ordered** - specifies that commands in the drive's task set are to be executed in the order received (FIFO).

The DQue bit (disable queuing) of the Control Mode page can be used to disable tagged command queuing.

6.11 COMMAND REORDERING

The drive uses Optimized Reordering Command Algorithm (ORCA) and tagged command queuing to provide greater throughput by reordering any commands queued on the drive to minimize rotational latency, seek time and head selection time. The performance improvement is significant in heavily loaded applications where queue depth exceeds 10-15 commands.

6.12 BANDED RECORDING

In banded recording, the disk is divided into multiple bands (also called partitions, notches, or bit-zoned areas). Starting at the inner band, each band further out has more blocks per track (a higher recording frequency). This use of multiple-frequency recording increases the capacity of the drive.

6.13 SPECIAL FUNCTIONS

Quantum has incorporated a number of unique options into the drive. The Quantum (Vendor) Special Function Control page is used to control these options. The options include:

- Initiate Synchronous Data Transfer Negotiation – a toggle-type parameter that allows the initiator to move between asynchronous and synchronous modes of data transfer.
- Write Protect – Prevents writing to the disk drive.
- Spin Delay – specifies if the drive should spin-up at power on or wait for a START STOP UNIT command from the host.
- Stagger Spin Delay – allows a series of drives to be spun up individually, reducing the load on the system power supply.

6.14 POWER MANAGEMENT

The Quantum Atlas 10K II Ultra 160/m SCSI hard disk drives have three basic power states:

- OFF – power is removed from the drive.
- UNIT READY – the drive is powered on and is ready to accept and execute commands. A single sub-state of UNIT READY exists, ACTIVE. This is the highest power consumption state of the drive. A media access command is acted upon immediately by the drive.

The drive remains in the UNIT READY state.

- NOT READY – The drive is powered on but cannot be addressed. An example of this state is during spin-up. In other cases the drive may require operator intervention. This state will return a Status = Check condition.

The Power Condition Page, controlled with the MODE SELECT command, sets the drive Power Management functions. **This page is not changeable.**

6.15 DIAGNOSTICS

The drive has extensive diagnostic capabilities, including those described below: power-on self-test, periodic self-adjustments, and host diagnostics. Quantum uses self-diagnostic tests and sense data tracking to manage drive errors, in addition to logical block address (LBA) revectoring.

6.15.1 Power On Self Test (POST)

The disk drive performs a self-diagnostic test immediately at power up. Both the Fault LED and a Busy LED briefly illuminate during spin-up as a lamp test.

If the drive passes the self-tests (and the spin-up jumper is installed), it spins up and performs further diagnostics on the Head Disk Assembly (HDA). The drive indicates “not ready” until the HDA tests are completed successfully.

The drive diagnostics test the following:

- RAM and ROM chips
- Buffer memory
- Gate array chips
- Read/write encoding circuits
- Basic servo functions

After a SELECTION TIME following power-on, the disk drive is able to respond with appropriate status and sense data to the TEST UNIT READY, INQUIRY, and REQUEST SENSE commands.

6.15.2 Periodic Self-Adjustments

During normal drive operation, certain periodic functions occur automatically so the drive can verify and maintain the integrity of its operations and optimize performance. The operations performed include those listed below.

- Move/Switch idle heads
- Adaptive seek profile calibration
- ROM checksum

Such periodic functions can make the drive appear to be active when no commands are outstanding, because the drive appears to be seeking. Periodic calibrations occur only after the drive has been inactive for 30 seconds.

NOTE
If a media access command is received during a calibration, the drive aborts the calibration process and acts on the command.

If desired, calibrations can be suspended to reduce drive overhead by setting the Manual Calibration bit in the (Vendor) Special Function Control page. If this is done, manual calibrations can still be initiated by performing a REZERO UNIT command.

6.15.3 Host Diagnostics

The Host Processor can command the drive to perform diagnostic tests and report status using RECEIVE and SEND DIAGNOSTIC RESULTS commands. Certain parameters can be sent to the drive to alter or enhance the performance of the diagnostic tests specified.

6.15.4 Log Sense

Quantum implements multiple LOG SENSE pages which logically group together parameters related to disk and subsystem activities. Monitoring this data enables users to observe drive performance and possibly predict failures.

6.16 ERROR RECOVERY

The mode pages control all error recovery. Any error recovery that is enabled is attempted without host intervention.

6.16.1 Seek Error Recovery

Seek error recovery is automatic if the hardware is physically able to position to the correct track.

6.16.2 Data Read Errors

Data read errors are recovered, if possible, by one or more of the following methods:

- Correction of the data “on-the-fly” with ECC hardware.
- Re-try the operation.
- Reload read channel registers to calibrated values.
- Off-track Recovery. Use different off-track offsets and re-try.
- Use a fixed gain while retrying.
- Margin servo thresholds to facilitate seek and detent.
- Margin the error tolerance of sync mark detection.
- Force sync mark detection while retrying.
- Skip Sector Rereads. Retries, each without one of eight earlier servo blocks.
- Correction of the bad data with the software ECC algorithm.
- Switching the bias current of the MR head between retries.
- TA baseline compensation, and other TA recovery methods.

6.16.3 Reallocation of Bad Blocks

Most of the spare blocks are located on the inner cylinders of the drive. All defective blocks are in-line spared at the factory or at format time. Some spare blocks are allocated at the end of each band, and are reserved for grown, revectoring blocks. All grown defects are reallocated to one of the spare blocks, using a “nearest neighbor” algorithm.

If bad block replacement is enabled (via the AWRE and ARRE bits in the Read-Write Error Recovery Page), blocks that are difficult to read or write may be replaced after multiple retries.

If the bad block data is *recoverable*, the drive first verifies that the original block location is bad with multiple write tests. If the write tests fail, the original data is reassigned to a new location.

If the block's data is *unrecoverable*, the block may or may not be replaced according to the state of the RUEE bit in the Quantum Vendor Unique Page. If the RUEE bit is set, the bad data is relocated with a Bad Data Mark appended to it. If the PER bit is set in the Read-Write Recovery Page, operations resulting in bad block replacement activate Check

6.17 THE ULTRA 160/m LOW VOLTAGE DIFFERENTIAL (LVD) SCSI INTERFACE

SCSI-3 provides increased performance and versatility to SCSI disks. LVD uses lower level voltage swings, differential signaling, and double transition (DT) clocking to allow a maximum bus speed of 160MB/second in wide SCSI configurations.

The Atlas 10K II disk drive supports the LVD/MSE (Multi-Mode LVD and Single-Ended) interface standard. When installed on a bus with only LVD devices the Atlas 10K disk drive operates in LVD mode at transfer rates up to 150 MB/second. It is fully compatible with Ultra2 LVD devices operating at slower transfer rates. If an Atlas 10K disk drive is installed on a bus with one or more single-ended (SE) devices, it automatically switches modes to operate in SE mode at the slower, legacy data rates.

Differential SCSI operation has the advantage of higher reliability of data transfer through increased immunity to electrical noise. Differential signaling uses a two-wire active system in which current from each wire flows in opposite directions and reverses direction for a signal transition. The direction of current flow determines whether the voltage levels on the two wires have a high/low or low/high relationship to each other. The common mode, or reference, level for the two LVD signals is +1.25 volts. High and low is defined relative to this level. LVD permits a wide range of signal amplitudes. The Atlas 10K interface drivers are designed to provide a high to low range on each signal of approximately 400 millivolts in a nominal configuration, resulting in a differential signal of approximately 800 millivolts.

Single-ended SCSI, in contrast, uses one-wire active signaling with the signal return wire connected to ground. The active signal range is switched between a high in the range of +2.4 – 3.0 volts and a low in the range of 0 - +.5 volts. LVD/MSE SCSI devices in single-ended mode do not have the signal return wires connected directly to ground, as these same pins must be driven to LVD levels when the device is operating in LVD mode. Instead, when in single-ended mode, the device turns on a driver transistor for each return line that acts as a switch to connect the signal return to ground.

6.17.1 DIFFSENS

The DIFFSENS signal in the SCSI bus defines the current operating mode of the SCSI bus. A level below +.5 volts defines a single-ended bus. A level between +.7 and +1.9 volts defines an LVD bus. A level above +2.4 volts defines a High Voltage Differential (HVD) bus. Before LVD was introduced, HVD was simply called Differential SCSI.

All single-ended-only devices connect the DIFFSENS signal to ground. On an LVD bus, the terminators drive the DIFFSENS signal to +1.3 volts unless some device is holding the signal at ground. HVD devices pull DIFFSENS up to +5 volts via a resistor.

The Atlas 10K disk drive is designed with a circuit that monitors the level of DIFFSENS at power-on. The drive then enables the appropriate circuits for LVD or single-ended operation, or disables its interface completely if it detects an HVD level on DIFFSENS.

The two normal operating environments for the Atlas 10K disk drive are an all-LVD or (LVD/MSE) bus, in which case it operates in its LVD mode, or a bus that contains at least one single-ended device, in which case the Atlas 10K disk drive operates in single-ended mode.

HVD mode is not supported by the Atlas 10K disk drive. Signal levels on an HVD bus can reach as high as +15 volts due to allowed common mode transients.

CAUTION

The LVD and HVD alternatives are mutually exclusive and damage may occur to SCSI devices if an LVD device is plugged into an operating HVD bus.


6.17.2 DIFFSENS SWITCHING

After establishing its initial operating mode, the Atlas 10K II disk drive continues to monitor the DIFFSENS signal. If a change in level is detected, such as might occur if a device is added to or removed from the bus, the DIFFSENS circuit detects the change and presents a mode change interrupt to the drive firmware. The operating mode of the interface switches to the new mode 100 msec after the change is detected, providing the new level has remained constant. This delay is required to prevent spurious mode switches due to noise on the DIFFSENS line.

6.17.3 DT AND ST CLOCKING

With 160/m interface, Quantum has introduced a new method of clocking data on the SCSI bus. This method, called Double Transition (DT) clocking, uses both the rising and falling edges of the REQ and ACK signals to clock data into the receiving device. This has the advantage of allowing the REQ and ACK signals to run at half the frequency that would be required under the traditional Single Transition (ST) SCSI clocking scheme for the same data rate. Lower frequency clocks result in a more reliable bus. Figure 6-1 illustrates the difference between DT and ST clocking.

Single Transition (ST) Clocking



Double Transition (DT) Clocking




Figure 6-1 *ST Clocking and DT Clocking*

The Atlas 10K II disk drive can operate at speeds up to 80 MB/sec in either ST or DT modes, depending on the capabilities of the host adapter. Transfer rates above 80 MB/sec are only supported in DT mode, as required by the SPI-3 SCSI standard.

6.17.4 CYCLIC REDUNDANCY CHECKING

When transferring data with DT clocking on the SCSI bus, a significantly improved form of error checking, cyclic redundancy checking (CRC) is used instead of the traditional parity checking. With CRC checking, the device sending data generates a four-byte CRC character based on the contents of the bytes being transferred. The receiving device also generates a CRC character as it receives the bytes. At the end of a data burst, the sending device transfers its CRC character and the receiving device compares the sender's CRC with its own. If the two CRC characters match, then the data burst transferred without error.

In ST mode, the Atlas 10K II disk drive uses the traditional parity checking method to ensure error free data transfers.

6.17.5 DOMAIN VALIDATION

A new feature in 160/m interface is the capability of the initiator and target devices to execute a test sequence of commands before customer operation begins to ensure that the SCSI bus can support the desired transfer rate. If errors are found, the initiator can negotiate different transfer parameters until a working configuration is found.

- Domain validation can detect configuration problems such as:
- Wide devices with a narrow bus segment between them
- DT devices connect through an ST expander
- Broken connections

It may also detect poor quality cables, marginal quality signals, and improper termination.

Domain validation is implemented using standard SCSI commands and new versions of the WRITE BUFFER and READ BUFFER commands. The sequence of commands used for domain validation may vary from system to system. The selection of commands used is entirely under control of host firmware and software.

6.18 HOT PLUGGING/REMOVAL AND INSERTION OF SCSI DEVICES

The Quantum Atlas 10K II SCA low voltage differential (LVD) disk drives are designed for use in "Hot Swap" applications within a properly designed and configured SCSI system. With care, system integrators can design storage arrays and/or SCSI buses using Quantum Atlas 10K ii disk drives that are "Hot Swap Case 4" tolerant.

Case 4 is defined as follows:

- Devices are powered and the bus may have active I/O processes ongoing, but the device being removed or inserted must be idle (no ongoing I/O processes during the insertion/removal).
- Ground connections to the drive must be made and maintained for 1 millisecond before, during, and 1 millisecond after the insertion/removal. The SCA-2 connector used on the Atlas 10K II disk drive meets this requirement.
- Device circuitry connected to bus pins must remain "glitch" free during power up or power down.

NOTE

In a multimode environment, any insertion or removal that changes the bus mode causes a transceiver mode change reset event.

System Considerations

It is not possible for the suppliers of various system components (e.g., targets, initiators, backplanes, terminators) to guarantee that a system can operate under “Hot Swap Case 4” conditions. Therefore, the system integrator bears the responsibility for ensuring that the system can meet “Hot Swap Case 4” operational criteria.

There are two metrics key to ensuring “Hot Swap Case 4” tolerance:

1. When a drive is being hot inserted into an active bus, a smaller capacitance on the SCSI bus pins create a smaller (that is, a lesser charge) transient spike on the bus.
2. When the drive is the closest drive to the connector at which another drive is being hot inserted, it is desirable that the input receivers ignore short low amplitude transients (usually via a low pass filter).

LVD SCSI devices may require more stringent system design to tolerate transients that occur during Case 4 insertion or removal. System integrators should consider the following when designing their backplanes or buses:

1. Larger connector-to-connector spacing causes the transients associated with hot swapping to be attenuated dramatically before reaching an operating drive.
2. Larger SCSI signal-trace capacitance (that is, lower impedance) is preferable for minimizing transient spikes, but can cause other bus problems unrelated to hot swapping. Draft standard SCSI-3 SCSI SPI-2 (SCSI Parallel Interface), Revision 18, Section 6.6.2 illustrates the tradeoffs between trace capacitance and drive-to-drive spacing.
3. The universe of drives needs to be homogeneous. A 15 pF drive hot-inserted into a rack of 15 pF drives presents no difficulty. A 30 pF drive inserted into a rack of 10 pF drives with very close connector spacing, however, is more likely to cause problems.

Appendix A

Quick Reference

Appendix A is a synopsis of SCSI commands, associated pages, messages, sense keys and status codes used in this manual.

A.1 SCSI-2 / SCSI-3 Equivalent Terminology

Table A-1 contains the SCSI-3 equivalent terms for common terms used in SCSI-2 documents. Table A-2 contains the SCSI-3 commands and op codes.

Table A-1 SCSI-2 / SCSI-3 Equivalent Terminology

SCSI-2	SCSI-3
abort	abort task set
abort tag	abort task
bus device reset	target reset
clear queue	clear task set
command complete	task complete
continue I/O process	continue task
head of queue tag	head of queue
ordered queue tag	ordered
incorrect initiator connection	overlapped commands
I/O process	task
phase	service
queue	task set
queue full	task set full
simple queue tag	simple

Table A-2 SCSI-3 Quick Reference – Commands and Messages

Command	Op Code (Hex)
CHANGE DEFINITION	40
FORMAT UNIT	04
INQUIRY	12
LOG SELECT	4C
LOG SENSE	4D
MODE SELECT (6)	15
MODE SELECT (10)	55
MODE SENSE (6)	1A
MODE SENSE (10)	5A
PERSISTENT RESERVE IN	5E
PERSISTENT RESERVE OUT	5F
READ (6)	08
READ (10)	28
READ BUFFER	3C
READ CAPACITY	25
READ DEFECT DATA	37
READ LONG	3E
READ MASK SKIP	E8
REASSIGN BLOCKS	07
RECEIVE DIAGNOSTIC RESULTS	1C
RELEASE (6)	17
RELEASE (10)	57
REPORT DEVICE IDENTIFIER	A3
REPORT LUNS	A0
REQUEST SENSE	03
RESERVE (6)	16
RESERVE (10)	56
REZERO UNIT	01
SEEK (6)	0B
SEEK (10)	2B
SEND DIAGNOSTIC	1D
SET DEVICE IDENTIFIER	A4
START/STOP UNIT	1B
SYNCHRONIZE CACHE	35
TEST UNIT READY	00
VERIFY	2F
WRITE (6)	0A
WRITE (10)	2A
WRITE AND VERIFY	2E
WRITE BUFFER	3B
WRITE LONG	3F
WRITE MASK SKIP	EA
WRITE SAME	41
XDREAD	52
XDWRITE	50
XPWRITE	51

(continued)

Table A-2 SCSI-3 Quick Reference – Commands and Messages (continued)

Messages	Msg. Code - Hex
ABORT	06
ABORT TAG	0D
BUS DEVICE RESET	0C
CLEAR QUEUE	0E
COMMAND COMPLETE	00
CONTINUE I/O PROCESS	12
DISCONNECT	04
HEAD OF QUEUE TAG	21
IDENTIFY	80 – FF
IGNORE WIDE RESIDUE	25
INITIATOR DETECTED ERROR	05
LINKED COMMAND COMPLETE	0A
LINKED COMMAND COMPLETE w/FLAG	0B
MESSAGE PARITY ERROR	09
MESSAGE REJECT	07
NO OPERATION	08
ORDERED QUEUE TAG	22
RELEASE RECOVERY	10
SAVE DATA POINTER	02
SIMPLE QUEUE TAG	20
SYNCHRONOUS DATA TRANSFER REQUEST	01
TARGET TRANSFER DISABLE	13
WIDE DATA TRANSFER REQ.	03

Table A-3 SCSI-3 Quick Reference - Pages

Pages	Page Code	Cmd
Application Client	0F	4D
ASCII Implemented Operating Definition	82	12
ASCII Information	01, 02, 03	12
Bad Block Replacement Summary	39	4D
Buffer Overruns and Underruns	01	4D
Caching	08 *	15, 55
Control Mode	0A *	15, 55
DER Description	3A	4D
Device Identification	83	12
Disk Error Recovery Counters	3A	4D
Disconnect-Reconnect	02	15, 55
Early Warning Status	36	4D
ECC Counters and Summary	3D	4D
Firmware Revision	C0	12
Format Device	03	1A
Format Status	08	4D
Implemented Operating Definition	81	12
Last n Error Events	07	4D
Medium Types Supported	0B	1A
Negotiated Rate Information	4C	12
Non-Medium Error Counter	06	4D
Notch and Partition	0C	15, 55
Peripheral Device	09	1A, 5A
Power Condition	0D	15
Quantum Manufacturing Number	C1	12
Quantum-Unique	39	15, 55
Read Error Counter	03	4D
Read-Write Error Recovery	01	15
Rigid Disk Geometry	04	15, 55
SCSI Bus Events	3E	4D
Seek Performance Summary	37	4D
Self Test Results	10	4D
Servo Events Counter	38	4D
Standard Inquiry Data	**	
Supported Diagnostic Pages	00	1C
Supported Log Pages	00	4D
Supported Vital Product Pages	00	12
Translate Address	40	1C
Unit Attention Control	00	15, 55
Unit Serial Number	80	12
Vendor Part Number	C1	12
Vendor-Unique Caching	38	1A, 5A
Verify Error Counter	05	4D
Verify Error Recovery	07	15, 55
Write Error Counter	02	4D
XOR Control	10	15, 55

* Not supported in SCSI-1/CSS mode

Table A-4 SCSI-3 Quick Reference – Sense Keys

Sense Keys	Sense Code - Hex
NO SENSE	0
RECOVERED ERROR	1
NOT READY	2
MEDIUM ERROR	3
HARDWARE ERROR	4
ILLEGAL REQUEST	5
UNIT ATTENTION	6
DATA PROTECT	7
VENDOR SPECIFIC	9
ABORTED COMMAND	B
MISCOMPARE	E

Table A-5 SCSI-3 Quick Reference – Status Codes

Status Codes	Status Code - Hex
GOOD	00
CHECK CONDITION	02
BUSY	08
RESERVATION CONFLICT	18
TASK SET FULL	28

Appendix B

SCSI REFERENCE INFORMATION

Appendix B provides information about SCSI bus timing, SCSI signal states, SCSI bus conditions, and SCSI message format and supported SCSI messages. Information in this appendix covers both SCSI-2 and SCSI-3 interfaces. Note that this information is general in nature, and is not specific to the ATLAS 10K II disk drive. It is provided as a reference only.

B.1 SCSI-2 Bus Timing

Table B-1 provides timing values and definitions for SCSI-2 SCSI bus signals.

Table B-1 SCSI Bus Timing Values -- ANSI SCSI-2 Standard

Timing Description	Value	Description
Arbitration Delay	2.4 μ s	Minimum time a SCSI device waits from asserting BSY for arbitration until the DATA BUS can be examined to see if arbitration has been won; there is no maximum time.
Assertion Period	90 ns	Minimum time a drive asserts REQ while using standard synchronous data transfers; also, the minimum time that an initiator asserts ACK while using standard (slow) synchronous data transfers.
Bus Clear Delay	800 ns	Maximum time for a SCSI device to stop driving all bus signals after: <ol style="list-style-type: none"> 1. BUS FREE is detected. 2. SEL is received from another SCSI device during ARBITRATION. 3. Transition of RST to true. For condition 1, the maximum time for a SCSI device to clear the bus is 1200 ns (1.2 μ s) from BSY and SEL first becoming both false. If a SCSI device requires more than a bus settle delay to detect BUS FREE, it clears the bus within a bus clear delay minus the excess time.
Bus Free Delay	800 ns	Maximum time a SCSI device waits from its detection of BUS FREE until its assertion of BSY when going to ARBITRATION.
Bus Set Delay	1.8 μ s	Maximum time for a device to assert BSY and its SCSI ID bit on the DATA BUS after it detects BUS FREE to enter ARBITRATION.
Bus Settle Delay	400 ns	Minimum time to wait for the bus to settle after changing certain control signals as defined in protocol definitions.

Table B-1 SCSI Bus Timing Values -- ANSI SCSI-2 Standard (continued)

Timing Description	Value	Description
Cable Skew Delay	10 ns	Maximum difference in propagation time allowed between any two SCSI bus signals measured between any two SCSI devices.
Data Release Delay	400 ns	Maximum time for an initiator to release the DATA BUS signals following the transition of the I/O signal from false to true.
Deskew Delay	45 ns	Minimum time required to wait for all signals (especially data signals) to stabilize at their correct, final value after changing.
Disconnection Delay	200 μ s	Minimum time that a drive waits after releasing BSY before participating in an ARBITRATION when honoring a DISCONNECT message from the initiator.
Hold Time	45 ns	Minimum time added between the assertion of REQ or ACK and changing the data lines to provide hold time in the initiator or drive while using standard (slow) synchronous data transfers.
Negation Period	90 ns	Minimum time that a drive negates REQ while using standard synchronous data transfers; also, the minimum time that an initiator negates ACK while using standard (slow) synchronous data transfers.
Power-On to Selection Time	10 s*	Recommended maximum time from power application until a drive is able to respond with appropriate status and sense data to the TEST UNIT READY, INQUIRY, and REQUEST SENSE commands.
Reset to Selection Time	250 ms*	Recommended maximum time after a hard RESET condition until a drive is able to respond with appropriate status and sense data to the TEST UNIT READY, INQUIRY, and REQUEST SENSE commands.
Reset Hold Time	25 μ s	Minimum time for which RST is asserted; there is no maximum time.
Selection Abort Time	200 μ s	Maximum time that a drive (or initiator) takes from its most recent detection of being selected (or reselected) until asserting a BSY response.
Selection Time-Out Delay	250 ms*	Recommended minimum time a SCSI device should wait for a BSY response during SELECTION or RESELECTION before starting the time-out procedure.
Transfer Period		Minimum time allowed between the leading edges of successive REQ pulses and of successive ACK pulses while using standard or fast synchronous data transfers. The period range is 200 to 500 ns minimum, standard, or 100 to 500 ns minimum, fast-synchronous. Note: negotiated and set during an SDTR message.

* Recommended time

(continued)

Table B-1 SCSI Bus Timing Values -- ANSI SCSI-2 Standard (continued)

Fast Synchronous Option		
Fast Assertion Period	30 ns	Minimum time a drive asserts REQ and the minimum time an initiator asserts ACK while using fast synchronous data transfers.
Fast Cable Skew Delay	5 ns	Maximum difference in propagation time allowed between any two SCSI bus signals measured between any two SCSI devices while using fast synchronous data transfers.
Fast Deskew Delay	20 ns	Minimum time required to wait for all signals (especially data signals) to stabilize at their correct, final value after changing while using fast synchronous data transfers.
Fast Hold Time	10 ns	Minimum time added between the assertion of REQ or ACK and the changing of the data lines to provide hold time in the initiator or drive, respectively while using fast synchronous data transfers.
Fast Negation Period	30 ns	Minimum time that a drive negates REQ and the minimum time that an initiator negates ACK while using synchronous data transfers.

B.2 SCSI-3 Bus Timing

SCSI devices that use a single transition (ST), i.e., they use only one edge of the REQ/ACK signal to clock data, can provide data transfer rates of up to 80 MB per second. Ultra 160/m SCSI-3, on the other hand, uses both edges of REQ/ACK, or double transition (DT) to double the speed of the bus.

Table B-2 provides the timing values for SCSI bus controls defined by the SCSI-3 standard. Note that the QAS (Quick Arbitrate/Select) items are not supported by the drive. Tables B-3 and B-4 provide SCSI bus data and information phase timing values for ST and DT devices, respectively.

Table B-2 SCSI-3 Bus Control Timing

Description	Values
Arbitration Delay	2.4 μ s
Bus Clear Delay	800 ns
Bus Free Delay	800 ns
Bus Set Delay	1.6 μ s
Bus Settle Delay	400 ns
Cable Skew (note 1)	4 ns
Data Release Delay	400 ns
DIFFSENS voltage filter time	100 ms
Physical Disconnection Delay	200 μ s
Power on to Selection (note 2)	10 s
QAS Arbitration Delay	1000 ns
QAS Assertion Delay	200 ns
QAS Release Delay	200 ns
QAS non-DATA phase REQ (ACK) period	50 ns
Reset Delay	200 ns
Reset Hold Time	25 μ s
Reset to Selection (note 2)	250 ms
Selection Abort Time	200 μ s
Selection Time-out Delay (note 2)	250 ms
System Deskew Delay	45 ns

Notes: 1 Cable Skew is measured at each device connection with the transmitted skew subtracted from the received skew.

2 This is a recommended time. It is not mandatory.

Table B-3 Single Transition (ST) SCSI Bus Data and Information Phase Timing

Description	Values (Note 4)				
	Asynch	Fast-5	Fast-10	Fast-20	Fast-40
ATN Transmit Setup Time	90 ns	33 ns	33 ns	21.5 ns	19.25ns
ATN Receive Setup Time	45 ns	17 ns	17 ns	8.5 ns	6.75 ns
Cable Skew (note 1)	4 ns	4 ns	4 ns	3 ns	2.5 ns
Receive Assertion Period (note 2)	N/A	70 ns	22 ns	11 ns	6.5 ns
Receive Hold Time (note 2 and note 3)	N/A	25 ns	25 ns	11.5 ns	4.75 ns
Receive Negation Period (note 2)	N/A	70 ns	22 ns	11 ns	6.5 ns
Receive Setup Time (note 2 and note 3)	N/A	15 ns	15 ns	6.5 ns	4.75 ns
Receive REQ (ACK) Period Tolerance	N/A	1.1 ns	1.1 ns	1.1 ns	1.1 ns
Signal Timing Skew	8 ns	8 ns	8 ns	5 ns	4.5 ns
REQ (ACK) Period	N/A	200 ns	100 ns	50 ns	25 ns
Transmit Assertion Period (note 2)	N/A	80 ns	30 ns	15 ns	8 ns
Transmit Hold Time (note 2 and note 3)	N/A	53 ns	33 ns	16.5 ns	9.25 ns
Transmit Negation Period (note 2)	N/A	80 ns	30 ns	15 ns	8 ns
Transmit Setup Time (note 2 and note 3)	N/A	23 ns	23 ns	11.5 ns	9.25 ns
Transmit REQ (ACK) Period Tolerance	N/A	1 ns	1 ns	1 ns	1 ns

- Notes:**
- 1 Cable Skew is measured at each device connection with the transmitted skew subtracted from the received skew.
 - 2 See SPI-3 specifications for measurement points.
 - 3 See SPI-3 specifications for examples of how to calculate setup and hold timing.
 - 4 SCSI bus timing values specified by the maximum transfer rate for the given range shall apply even if a slower transfer rate within the given range is negotiated.

Table B-4 Double Transition (DT) SCSI Bus Data and Information Phase Timing

Timing Description	Values (Note 4)			
	Fast-10	Fast-20	Fast-40	Fast-80
ATN Transmit Setup Time	48.4 ns	29.2 ns	19.6 ns	14.8 ns
ATN Receive Setup Time	13.6 ns	7.8 ns	4.9 ns	3.45 ns
Cable Skew (note 1)	4 ns	3 ns	2.5 ns	2.5 ns
pCRC Receive Hold Time	10.2 ns	5.1 ns	2.55 ns	1.45 ns
pCRC Receive Setup Time	20.2 ns	15.1 ns	12.55 ns	11.45 ns
pCRC Transmit Hold Time	37 ns	18.5 ns	9.25 ns	4.8 ns
pCRC Transmit Setup Time	47 ns	28.5 ns	19.25 ns	14.8 ns
Receive Assertion Period (note 2)	80 ns	40 ns	20 ns	10 ns
Receive Hold Time (note 2 and note 3)	11.6 ns	5.8 ns	2.9 ns	1.45 ns
Receive Negation Period (note 2)	80 ns	40 ns	20 ns	10 ns
Receive Setup Time (note 2 and note 3)	11.6 ns	5.8 ns	2.9 ns	1.45 ns
Receive REQ (ACK) Period Tolerance	0.7 ns	0.7 ns	0.7 ns	0.7 ns
Signal Timing Skew	26.8 ns	13.4 ns	6.7 ns	3.35 ns
REQ (ACK) Period	200 ns	100 ns	50 ns	25 ns
Transmit Assertion Period (note 2)	92 ns	46 ns	23 ns	11.5 ns
Transmit Hold Time (note 2 and note 3)	38.4 ns	19.2 ns	9.6 ns	4.8 ns
Transmit Negation Period (note 2)	92 ns	46 ns	23 ns	11.5 ns
Transmit Setup Time (note 2 and note 3)	38.4 ns	19.2 ns	9.6 ns	4.8 ns
Transmit REQ (ACK) Period Tolerance	0.6 ns	0.6 ns	0.6 ns	0.6 ns

- Notes:**
- 1 Cable Skew is measured at each device connection with the transmitted skew subtracted from the received skew.
 - 2 See SPI-3 specifications for measurement points.
 - 3 See SPI-3 specifications for examples of how to calculate setup and hold timing.
 - 4 SCSI bus timing values specified by the maximum transfer rate for the given range shall apply even if a slower transfer rate within the given range is negotiated.

The Timing Parameters and Descriptions for SCSI-3 ST/DT are defined in alphabetical order below:

Arbitration delay - The minimum time a SCSI device shall wait from asserting the BSY signal for arbitration until the DATA BUS is examined to see if arbitration has been won. There is no maximum time.

ATN Transmit Setup Time - The minimum time provided by the transmitter between the assertion of the ATN signal and the negation of the ACK signal.

ATN Receive Setup Time - The minimum time required at the receiver between the assertion of the ATN signal and the negation of the ACK signal to recognize the assertion of an Attention Condition.

Bus Clear Delay - The maximum time for a SCSI device to release all SCSI bus signals after:

- a) the BUS FREE phase is detected (the BSY and SEL signals are both false for a bus settle delay);
- b) the SEL signal is received from another SCSI device during the ARBITRATION phase;
- c) the transition of the RST signal to true.

For item a) above, the maximum time for a SCSI device to release all SCSI bus signals is 1200 ns from the BSY and SEL signals first becoming both false. If a SCSI device requires more than a bus settle delay to detect BUS FREE phase, it shall release all SCSI bus signals within a bus clear delay minus the excess time.

Bus Free Delay - The minimum time that a SCSI device shall wait from its detection of the BUS FREE phase (BSY and SEL both false for a bus settle delay) until its assertion of the BSY signal in preparation for entering the ARBITRATION phase.

Bus Set Delay - The maximum time for a SCSI device to assert the BSY signal and its SCSI ID after it detects a BUS FREE phase for the purpose of entering the ARBITRATION phase.

Bus Settle Delay - The minimum time to wait for the bus to settle after changing certain control signals as called out in the protocol definitions.

Cable Skew - The maximum difference in propagation time allowed between any two SCSI bus signals measured between any two SCSI devices excluding any signal distortion skew delays.

pCRC Receive Hold Time - The minimum time required at the receiver between the transition of the REQ signal and the transition of the P_CRCA signal while pCRC protection is enabled.

pCRC Receive Setup Time - The minimum time required at the receiver between the transition of the P_CRCA signal and the transition of the REQ signal while pCRC protection is enabled.

pCRC Transmit Hold Time - The minimum time provided by the transmitter between the transition of the REQ signal and the transition of the P_CRCA signal while pCRC protection is enabled.

Receive Hold Time - For ST data transfers the minimum time required at the receiving SCSI device between the assertion of the REQ signal or the ACK signals and the changing of the DATA BUS while using synchronous data transfers.

For DT data transfers the minimum time required at the receiving SCSI device between the transition (i.e. assertion or negation) of the REQ signal or the ACK signals and the changing of the DATA BUS while using synchronous data transfers.

Receive Negation Period - The minimum time required at a SCSI device receiving a REQ signal for the signal to be negated while using synchronous data transfers.

Also, the minimum time required at a SCSI device receiving an ACK signal for the signal to be asserted while using synchronous data transfers.

For SE fast-5 and fast-10 operation, the time period is measured at the 2.0 V level. For SE fast-20 operation the period is measured at the 1.9 V level.

Receive Setup Time - For ST data transfers the minimum time required at the receiving SCSI device between the changing of DATA BUS and the assertion of the REQ signal or the ACK signal while using synchronous data transfers.

For DT data transfers the minimum time required at the receiving SCSI device between the changing of DATA BUS and the transition of the REQ signal or the ACK signal while using synchronous data transfers.

Receive REQ (ACK) Period Tolerance - The minimum tolerance that a SCSI device shall allow to be subtracted from the REQ (ACK) period.

REQ (ACK) Period - The REQ (ACK) period during synchronous data transfers is measured from an assertion edge of the REQ (ACK) signal to the next assertion edge of the signal.

In DT DATA phases the nominal transfer period for data is half that of the REQ (ACK) period during synchronous data transfers since data is qualified on both the assertion and negation edges of the REQ (ACK) signal.

In ST DATA phases the nominal transfer period for data is equal to the REQ (ACK) period during synchronous data transfers since data is only qualified the assertion edge of the REQ (ACK) signal.

Reset Delay - The minimum time that the RST signal shall be continuously true before the SCSI device shall initiate a reset.

Reset Hold Time - The minimum time that the RST signal is asserted. There is no maximum time.

Reset To Selection - The recommended maximum time from after a reset condition until a SCSI target is able to respond with appropriate status and sense data to the TEST UNIT READY, INQUIRY, and REQUEST SENSE commands (See SCSI Primary Commands-2 Standard).

Selection Abort Time - The maximum time that a SCSI device shall take from its most recent detection of being selected or reselected until asserting the BSY signal in response. This time-out is required to ensure that a target or initiator does not assert the BSY signal after a SELECTION or RESELECTION phase has been aborted.

B.3 Signal States

The following paragraphs describe the SCSI signal values and SCSI ID bits.

B.3.1 Signal Values

All signal values are actively driven true (low voltage). Because the signal drivers are OR-tied, the bias circuitry on the bus terminator pulls false when it is released by the drivers at every SCSI device. If any device asserts a signal, (e.g., OR-tied signals), the signal is true. Table B-5 shows the ANSI-specified and defined signal sources. Any device can assert RST at any time.

Table B-5 Signal Sources

SIGNALS						
Bus Phase	BSY	SEL	C/D I/O MSG REQ	ACK ATN	DB(7-0) DB(P)	DB(16-8) DB(P1)
BUS FREE	None	None	None	None	None	None
ARBITRATION	All	Winner	None	None	S ID	S ID
SELECTION	I&T	Init	None	Init	Init	Init
RESELECTION	I&T	Targ	Targ	Init	Targ	Targ
COMMAND	Targ	None	Targ	Init	Init	None
DATA IN	Targ	None	Targ	Init	Targ	Targ
DATA OUT	Targ	None	Targ	Init	Init	Init
STATUS	Targ	None	Targ	Init	Targ	None
MESSAGE IN	Targ	None	Targ	Init	Targ	None
MESSAGE OUT	Targ	None	Targ	Init	Init	None

All: The signal is driven by all SCSI devices that are actively arbitrating.

S ID: Each SCSI device that is actively arbitrating asserts its unique SCSI ID bit. The other seven (or fifteen) data bits are released. The parity bit, DB (P) or DB (P1), can be released or driven true, but is never driven false during this phase.

I&T: The signal is driven by the initiator, target (drive), or both, as specified in the SELECTION and RESELECTION phase.

Init: If driven, this signal is driven only by the active initiator.

None: The signal is released; that is, not driven by any SCSI device. The bias circuitry of the bus terminators pulls the signal to the false state.

Winner: The signal is driven by the winning SCSI device.

Targ: If the signal is driven, it is driven only by the active drive.

B.3.2 SCSI ID Bits & Priorities

The SCSI standard permits a maximum of eight SCSI devices on a narrow (8 address bits) SCSI bus and up to 16 devices on a wide (16 address bits) bus. Each SCSI device has a single, unique SCSI ID address bit assigned to it. This SCSI ID identifies and also sets the priority of the device on the bus. The bit priority is used for arbitration as well as the SCSI ID.

See Table B-6 for the SCSI ID data bus bit assignments and the respective bit priorities. Note that bit DB(7) has the highest priority for either 8-bit or wide (16-bit) transfer modes. On the drive, the SCSI ID is assigned by configuring jumpers or remote switches to the option connector in normal 3-bit (or 4-bit) binary fashion. Chapter 2 has full instructions on setting the SCSI ID jumpers.

Since the drives are SCAM-1 compliant, a SCAM initiator may assign a SCAM SCSI ID that is different than the one set with jumpers or switches.

Table B-6 SCSI ID Bits & Arbitration Priorities

DATA BUS	SCSI ID & PRIORITY (8-BIT TRANSFERS)		SCSI ID & PRIORITY (16-BIT TRANSFERS)	
	BIT NUMBER ¹	ID (ADDRESS)	PRIORITY	ID (ADDRESS)
DB(7)	7	1	7	1
DB(6)	6	2	6	2
DB(5)	5	3	5	3
DB(4)	4	4	4	4
DB(3)	3	5	3	5
DB(2)	2	6	2	6
DB(1)	1	7	1	7
DB(0)	0	8	0	8
DB(15)	-	-	15	9
DB(14)	-	-	14	10
DB(13)	-	-	13	11
DB(12)	-	-	12	12
DB(11)	-	-	11	13
DB(10)	-	-	10	14
DB(9)	-	-	9	15
DB(8)	-	-	8	16

¹ **NOTE:** A single bit is used to identify a device and define its priority on the SCSI bus. Highest priority (1) is at the top.

B.4 SCSI Signals

Table B-7 defines the SCSI-3 bus signals.

Table B-7 SCSI-3 Bus Signal Definitions

Signal	Definition
ACK (acknowledge)	A signal driven by the initiator as an acknowledgment of receipt of data from a target or as a signal to a target indicating when the target should read the data (out) lines.
ATN (attention)	A signal driven by an initiator to indicate that it has a message to send.
BSY (busy)	An OR-tied signal that indicates that the bus is in use.
C/D (control/data)	A signal driven by a target that indicates whether CONTROL or DATA information is on the DATA BUS. True (low voltage) indicates CONTROL.
DB(7–0,P) (data bus)	Eight data-bit signals, plus a parity-bit signal that form a DATA BUS. DB(7) is the most significant bit and has the highest priority (8 or 16-bit) during ARBITRATION. Bit number, significance, and priority decrease downward to DB(0). A data bit is defined as 1 when the signal value is true (low voltage) and 0 when the signal value is false (high voltage). Data parity DB(P) is odd. Parity is undefined during ARBITRATION.
DB(15–8,P1) (data bus)	Eight data-bit signals, plus one parity-bit signal, that forms an extension to the DATA BUS. They are used for 16-bit (wide) interfaces. DB(15) is the most significant bit and has the higher priority (but below bit DB(0) during ARBITRATION. Bit number, significance, and priority decrease downward to DB(8). Data Parity DB (P1) is odd.
I/O (input/output)	A signal driven by a target that controls the direction of data movement on the DATA BUS with respect to an initiator. True indicates input to the initiator. Also used to distinguish between SELECTION and RESELECTION modes.
MSG (message)	A signal driven by a target during the MESSAGE phase.
REQ (request)	A signal driven by a target to indicate a request for an information transfer to or from the initiator. Each byte of data transferred is accompanied with a REQ/ACK "handshake". See also, ACK.
RST (reset)	An OR-tied signal that initiates a RESET condition.
SEL (select)	An OR-tied signal used by an initiator to select a target or by a target to reselect an initiator.

B.5 SCSI Bus Phases

The SCSI communication architecture includes eight distinct phases:

- BUS FREE phase
- ARBITRATION phase
- SELECTION phase
- RESELECTION phase
- COMMAND phase
- DATA (In/Out) phases
- STATUS phase
- MESSAGE (In/Out) phases

The last four phases are called the “information transfer phases.”

The SCSI bus can never be in more than one phase at any given time. In the following descriptions, signals that are not mentioned are not asserted.

B.5.1 BUS FREE Phase

The BUS FREE phase indicates that there is no current I/O process and that the SCSI bus is available for a connection. SCSI devices detect the BUS FREE phase after the SEL and BSY signals are both false for at least one bus settle delay.

During normal operation, the BUS FREE phase is entered when the drive releases the BSY signal. However, the BUS FREE phase can be entered following the release of the SEL signal after a SELECTION or RESELECTION phase time-out.

The BUS FREE phase normally does not begin because of the drive’s release of the BSY signal unless it has occurred after the detection of a reset condition or after a drive has successfully transmitted or received one of the following messages:

Messages Transmitted from Drive:

- DISCONNECT
- COMMAND COMPLETE

Messages Received by Drive:

- ABORT TASK
- ABORT TASK SET
- TARGET RESET
- RELEASE RECOVERY
- CLEAR QUEUE

If an initiator detects the release of the BSY signal by the drive at any other time, the drive is indicating an error condition to the initiator. The drive can perform this transition to the BUS FREE phase independently of the state of the ATN signal. The initiator manages this condition as an unsuccessful I/O process termination. The drive terminates the I/O process by clearing all pending data and status information

for the affected nexus. The drive can optionally prepare sense data that can be retrieved by a REQUEST SENSE Command.

Bus Free Sequence

1. BSY and SEL signals are continuously false for one bus settle delay.

2. SCSI devices release all SCSI bus signals within one bus clear delay.

If a SCSI device requires more than one bus settle delay to detect the BUS FREE phase, then it releases all SCSI bus signals within one bus clear delay minus the excess time to detect the BUS FREE phase.

The total time to clear the SCSI bus cannot exceed one bus settle delay plus one bus clear delay.

B.5.2 ARBITRATION Phase

The ARBITRATION phase allows one SCSI device to gain control of the SCSI bus so that it can initiate or resume an I/O process.

A SCSI device arbitrates for the SCSI bus by asserting both the BSY signal and its own SCSI ID after a BUS FREE phase occurs. If a higher priority SCSI ID bit is asserted on the bus, the lower-priority SCSI device loses the arbitration. Narrow data bus modules/drives (8-bit) recognize SCSI IDs 7 through 0 in descending priority in arbitration. Wide data bus modules/drives (16-bit) recognize SCSI IDs 15 through 0 in descending order (ID bit numbers 7 through 0, then 15 through 8).

Arbitration Sequence

1. The SCSI device waits for the BUS FREE phase to occur.
2. The SCSI device waits a minimum of one bus free delay after detection of the BUS FREE phase before driving any signal.
3. The SCSI device arbitrates for the SCSI bus by asserting the BSY signal and its SCSI ID.
4. The SCSI device waits at least an arbitration delay to determine arbitration results:

NOTE

Step 4 requires that every device complete the arbitration phase to the point of SEL being asserted (for a SELECTION or RESELECTION phase) to avoid hanging the bus.

- If a higher priority SCSI ID bit is true on the DATA BUS, the SCSI device loses the arbitration.
- The losing SCSI device releases the BSY signal and its SCSI ID bit within one bus clear delay after the SEL signal asserted by the arbitration winner becomes true.
- The losing SCSI device waits for the SEL signal to become true before releasing the BSY signal and SCSI ID bit. The losing SCSI device returns to step 1
- If no higher priority SCSI ID bit is true on the DATA BUS, the SCSI device wins the arbitration and asserts the SEL signal.
- The winning SCSI device waits at least one bus clear delay plus one bus settle delay after asserting the SEL signal before changing any signals.

B.5.3 SELECTION Phase

The SELECTION phase allows an initiator to select a drive to initiate a drive function.

The SCSI device that won the arbitration has both the BSY and SEL signals asserted and has delayed at least one bus clear delay plus one bus settle delay before ending the ARBITRATION phase. The SCSI device that won the arbitration becomes an initiator by not asserting the I/O signal.

During SELECTION, the I/O signal is negated so that this phase can be distinguished from the RESELECTION phase.

Selection Sequence

The initiator:

1. Sets the DATA BUS to the OR of its SCSI ID bit and the drive's SCSI ID bit.
2. Asserts the ATN signal (signaling that a MESSAGE OUT phase is to follow the SELECTION phase).
3. Waits at least two deskew delays.
4. Releases the BSY signal.
5. Waits at least one bus settle delay.
6. Looks for a response from the drive.
7. No less than two deskew delays after the initiator detects that the BSY signal has been asserted by the target/drive, it releases the SEL signal and can change the DATA BUS.

The drive:

8. Determines that it is selected when the SEL signal and its SCSI ID bit are true and the BSY and I/O signals are false for at least one bus settle delay.
9. Can examine the DATA BUS to determine the SCSI ID of the selecting initiator.
10. Asserts the BSY signal within a selection abort time of its most recent detection of being selected.

This is required for correct operation of the selection time-out procedure.

The drive does not respond to a selection if bad parity is detected. Also, if more than two SCSI ID bits are on the DATA BUS, the drive does not respond to selection.

11. Waits until the SEL signal is false.
12. Asserts the REQ signal to enter an information transfer phase.

Selection Time-Out

Two optional time-out procedures are specified for clearing the SCSI bus if the initiator waits a minimum of a selection time-out delay and there has been no BSY signal response from the drive.

1. The initiator asserts the RST signal.
2. The initiator follows these steps:
 - a. continues asserting the SEL and ATN signals and releases the DATA BUS.
 - b. If it has not detected the BSY signal to be true after at least a selection abort time plus two deskew delays, the initiator releases the SEL and ATN signals, allowing the SCSI bus to go to the BUS FREE phase.

When responding to selection, SCSI devices ensure that the selection was still valid within a selection abort time of their assertion of the BSY signal. Failure to comply with the requirement could result in an improper selection.

B.5.4 RESELECTION Phase

RESELECTION is an optional phase that allows a drive to reconnect to an initiator to continue an operation that was previously started by the initiator but was suspended by the drive.

The initiator determines that it is reselected when the SEL and I/O signals and its SCSI ID bit are true, and the BSY signal is false for at least one bus settle delay.

Reselection Sequence

The drive:

1. Upon completing the ARBITRATION phase, asserts both the BSY and SEL signals.
2. Delays at least one bus clear delay plus one bus settle delay.
3. Asserts the I/O signal.

4. Sets the DATA BUS to the logical OR of its SCSI ID bit and the initiator's SCSI ID bit.
5. Waits at least two deskew delays.
6. Releases the BSY signal.
7. Waits at least one bus settle delay before looking for a response from the initiator.

The initiator:

8. Determines that it is selected when the following occur for at least one bus settle delay: SEL, I/O, and the initiator's SCSI ID bit are true and BSY is false.
9. Examines the DATA BUS to determine the SCSI ID of the reselecting drive.
10. Asserts the BSY signal within a selection abort time of its most recent detection of being reselected.

The initiator does not respond to a RESELECTION phase if bad parity is detected or if more than two SCSI ID bits are on the DATA BUS.

The drive:

11. Detects the BSY signal is true.
12. Asserts the BSY signal.
13. Waits at least two deskew delays.
14. Releases the SEL signal.

The drive can then change the I/O signal and the DATA BUS.

The initiator:

15. Detects the SEL signal is false.
16. Releases the BSY signal.

The drive:

17. continues asserting the BSY signal until it relinquishes the SCSI bus.

Reselection Time-Out

The Atlas 10K II drive performs the following recovery procedure in order to clear the SCSI bus in response to a reselection time-out. The drive waits a minimum of a selection time-out delay before initiating this procedure.

1. The drive continues asserting the SEL and I/O signals and releases the data bus signals.
2. The drive waits a for a total of a selection abort time plus two deskew delays.
3. The drive tests BSY. If it is asserted, the drive completes the reselection protocol.
4. If BSY is not asserted, the drive releases the SCSI bus.

B.5.5 Information Transfer Phases

The COMMAND, DATA, STATUS, and MESSAGE phases are known as the Information Transfer Phases because they are used to transfer data or control information.

The C/D, I/O, and MSG signals are used to distinguish between the different information transfer phases. The drive asserts these three signals and so controls all information transfer phase changes. The drive can also cause a BUS FREE phase by releasing the MSG, C/D, I/O, and BSY signals. The initiator can request a MESSAGE OUT phase by asserting the ATN signal.

The information transfer phases use one or more REQ/ACK handshakes to control the information transfer. Each REQ/ACK handshake allows the transfer of one byte of information. During the information transfer phases, the BSY signal remains true and the SEL signal remains false. Additionally, the drive continuously envelopes the REQ/ACK handshake(s) with the C/D, I/O, and MSG signals in such a manner that these control signals are valid for one bus settle delay before the assertion of the REQ signal of the first handshake and remain valid after the negation of the ACK signal at the end of the handshake of the last transfer of the phase.

After the negation of the ACK signal of the last transfer of the phase, the drive can prepare for a new phase by asserting or negating the C/D, I/O, and MSG signals. These signals can be changed together or individually. They can be changed in any order and can be changed more than once (although each line should change only once). A new phase does not begin until the REQ signal is asserted for the first byte of the new phase.

A phase ends when the C/D, I/O, or MSG signal changes after the negation of the ACK signal. The time between the end of a phase and the assertion of the REQ signal beginning a new phase is undefined. An initiator is allowed to anticipate a new phase based on the previous phase, the expected new phase, and early information provided by changes in the C/D, I/O, and MSG signals. However, the anticipated phase is not valid until the REQ signal is asserted at the beginning of the next phase.

Direction of Transfer

The direction of an information transfer is indicated by the I/O signal, as outlined in Table B-8. A True I/O Signal (1) indicates a transfer from the drive to the initiator. A False I/O Signal (0) indicates a transfer from the initiator to the drive.

Table B-8 Information Transfer Phases

Signal			Phase Name	Direction of Transfer/ Definition
MSG	C/D	I/O		
0	0	0	DATA OUT	Initiator to Drive. Allows the drive to request that data be sent from the initiator to the drive.
0	0	1	DATA IN	Drive to initiator. Allows the drive to send data to the initiator.
0	1	0	COMMAND	Initiator to Drive. Allows the drive to request a command from the initiator.
0	1	1	STATUS	Drive to initiator. Allows the drive to request that status information be sent from the drive to the initiator.
1	1	0	MESSAGE OUT	Initiator to Drive. Allows the drive to request that message(s) be sent from the initiator to the drive; the drive invokes this phase in response to the attention condition created by the initiator. The drive handshakes byte(s) until the ATN signal is negated, except when rejecting a message. <i>See Message Out — Additional Conditions.</i>
1	1	1	MESSAGE IN	Drive to initiator. Allows the drive to send message(s) to the initiator.

B.5.5.1 Asynchronous Data Transfer

Drive to Initiator Transfer Procedure

1. The drive drives the DB (7-0, P) signals to their desired values.
2. The drive delays at least one Deskew delay plus a cable skew delay.
3. The drive asserts the REQ signal.
4. The initiator reads the DB (7-0, P) signals.
5. The initiator indicates its acceptance of the data by asserting the ACK signal.
6. When ACK is true at the drive, drive can change or release the DB (7-0, P) signals.
7. Drive negates the REQ signal.
8. Initiator negates the ACK signal.
9. Drive can continue the transfer by driving the DB (7-0, P) signals and asserting the REQ signal (steps 1- 3).

Initiator-to-Drive Transfer Procedure

1. The drive asserts the REQ signal.
2. The initiator drives the DB (7-0, P) signals to their desired values.
3. The initiator delays at least one Deskew delay plus a cable skew delay.
4. The initiator asserts the ACK signal.
5. When ACK is true at the drive, drive reads the DB (7-0, P) signals.
6. The drive negates the REQ signal.
7. The initiator can change or release the DB (7-0, P) signals.
8. The initiator negates the ACK signal.
9. The drive can continue the transfer by asserting the REQ signal (step 1).

B.5.5.2 Synchronous Data Transfer

Synchronous Data Transfer is optional and is used only in DATA phases and only if a synchronous data transfer agreement is established. The REQ/ACK offset specifies the maximum number of REQ pulses that can be sent by the drive in advance of the number of ACK pulses received from the initiator, establishing a pacing mechanism. If the number of REQ pulses exceeds the number of ACK pulses set by the REQ/ACK offset, the drive does not assert the REQ signal until after the leading edge of the next ACK pulse is received. For successful completion of the data phase, the number of ACK and REQ pulses must be equal.

The initiator sends one ACK signal pulse for each REQ pulse received. The ACK signal can be asserted as soon as the leading edge of the corresponding REQ pulse has been received.

Drive-to-Initiator Transfer Procedure

1. The drive sets the DB (7-0, P) signals to the desired values.
2. The drive delays at least one Deskew delay plus a Cable Skew delay before asserting REQ.
3. The DB (7-0, P) signals are held valid for a minimum of one Deskew delay plus one Cable Skew delay plus one Hold Period after REQ is asserted.
4. The drive asserts the REQ signal for a minimum of one Assertion period after which the drive can negate the REQ signal.
5. After a Skew plus a Deskew plus a Hold Period delay, the drive can change or release the DB (7-0, P) signals.
6. The initiator reads the DB(7-0, P) signals within one Hold Period of the transition of the REQ signal to true.
7. The initiator indicates its acceptance of the data by asserting an ACK pulse.
8. The drive waits at least the greater of the following periods before asserting REQ again: a.) A transfer period from the last transition of the REQ signal to true. b.) A negation period from the last transition of the REQ signal to false.

9. The initiator waits at least the greater of the following periods before reasserting ACK, (assuming that REQ/ACK offset is set to zero): a.) A transfer period from the last transition of the ACK signal to true; b.) A negation period from the last transition of the ACK signal to false.
10. Otherwise the initiator can wait for an undefined period beyond the last received REQ pulse (or after the maximum, negotiated, offset has been reached) before asserting ACK. Note that when the maximum offset is reached, the target (drive) does not send any more REQ pulses until it receives at least one ACK pulse (since the REQ pulses must never be ahead of the ACK pulses by more than the offset).
11. Since the number of REQ and ACK pulses must be equal at the end of the phase, the target does not change the phase until all required ACK pulses have been received.

Initiator-to-Drive Transfer Procedure

1. Initiator transfers one byte for each REQ pulse received.
2. Drive asserts from one to fifteen REQ signals.
3. After receiving the leading edge of a REQ signal, initiator drives the DB (7-0, P) signals to the desired values (if no offset is to be applied to the transfer). If a REQ/ACK offset is used, the initiator waits up to the maximum, negotiated offset before driving the data lines.
The DB (7-0, P) signals are held valid for at least one deskew delay plus one Cable Skew Delay plus one Hold Period after the assertion of the ACK signal.
4. The initiator delays at least one Deskew Delay plus a Cable Skew Delay before asserting an ACK pulse.
5. After a Skew plus a Deskew plus a Hold Period delay, the initiator can change or release the DB (7-0, P) signals.
6. The initiator asserts the ACK signal for a minimum of one assertion period after which the initiator can negate the ACK signal.
7. The Drive reads the DB (7-0, P) signals within one Hold period of ACK signal going true.
8. The drive waits at least the greater of the following periods before again asserting the REQ signal:
 - a. A transfer period from the last transition of the REQ signal to true.
 - b. A negation period from the last transition of the REQ signal to false.
9. The initiator waits at least the greater of the following periods before again asserting the ACK signal unless an REQ/ACK offset greater than zero is used:
 - a. A transfer period from the last transition of the ACK signal to true.
 - b. A negation period from the last transition of the ACK signal to false.

B.5.5.3 Wide Data Transfer

Wide Data Transfer is optional and can be used in the DATA phase only if a wide transfer agreement is established. Wide data transfer is established by negotiations between the initiator and drive using the Wide Data Transfer Request message. The IWN bit in the Quantum (Vendor) Special Function Control page specifies which device (drive or initiator) is to initiate the negotiations.

During 16-bit wide data transfers, the first logical data byte for each data phase is transferred across the DB (0–7, P) signals and the second logical data byte is transferred across the DB (15–8, P1) signals. Subsequent pairs of data bytes are likewise transferred in parallel.

B.5.5.4 Signal Restrictions Between Phases

When the SCSI bus is between two information transfer phases, the following restrictions apply to the SCSI bus signals:

- The BSY, SEL, REQ, and ACK, signals do not change.
- The C/D, I/O, MSG, and DATA BUS signals can change.
- When changing the DATA BUS direction from out (initiator-driving) to in (drive-driving), the drive delays driving the DATA BUS by at least a data release delay plus one bus settle delay after asserting the I/O signal and the initiator releases the DATA BUS no later than a data release delay after the transition of the I/O signal to true.
- When switching the DATA BUS from in to out, the drive releases the DATA BUS no later than a deskew delay after negating the I/O signal.
- The ATN and RST signals can change as defined under the descriptions for the attention condition (see B.5.1) and reset condition (see B.5.2).

B.5.5.5 Message Out — Additional Conditions

If the drive detects one or more parity error(s) received, it can indicate its desire to retry the message(s) by asserting the REQ signal after detecting the ATN signal has gone false and prior to changing to any other phase. The initiator, upon detecting this condition, re-sends all the previous message byte(s) in the same order as previously sent during this phase. When re-sending more than one message byte, the initiator asserts the ATN signal at least two deskew delays prior to asserting the ACK signal on the first byte and maintains the ATN signal until the last byte is sent.

If the drive does not retry the MESSAGE OUT phase or it exhausts its retry limit, it can:

- Return CHECK CONDITION status and set the sense key to ABORTED COMMAND and the additional sense code to MESSAGE ERROR.
- Indicate an exception condition by performing an unexpected disconnect.

The drive can act on messages as received as long as no parity error is detected and can ignore all remaining messages sent under one ATN condition after a parity error is detected. When a sequence of messages is re-sent by an initiator because of a drive detected parity error, the drive does not act on any message that it acted on the first time the messages are received.

If the drive receives all the message byte(s) successfully, it indicates that it does not want to retry by changing to any information transfer phase other than the MESSAGE OUT phase and transfers at least one byte. The drive can also indicate that it has successfully received the message byte(s) by changing to the BUS FREE phase.

B.6 SCSI Bus Conditions

The SCSI bus has two asynchronous conditions: Attention and Reset.

B.6.1 Attention Condition

The attention condition informs a drive that an initiator has a message ready. The drive gets the message by performing a MESSAGE OUT phase. The attention condition requires the following timing:

- The initiator creates the attention condition by asserting ATN at any time except during the ARBITRATION or BUS FREE phases.
- The initiator negates the ATN signal at least two deskew delays before asserting the ACK signal while transferring the last byte of the message.

If the drive detects that the initiator failed to meet this requirement, then the drive goes to BUS FREE.

- Before transition to a new bus phase, the initiator asserts the ATN signal, then waits at least two deskew delays before negating the ACK signal for the last byte transferred in the current bus phase.
- Asserting the ATN signal later cannot be honored until a later bus phase and then cannot result in the expected action.

A drive responds with a MESSAGE OUT phase as described in Table B-9.

Table B-9 Drive MESSAGE OUT Phase Response

ATN Signal True in Phase...	The Drive Enters MESSAGE OUT...
COMMAND	After transferring part or all of the command descriptor block bytes.
DATA	At the drive's earliest convenience (often on a logical block boundary). The initiator continues REQ/ACK handshakes until it detects the phase change.
STATUS	After the status byte has been acknowledged by the initiator.
MESSAGE IN	Before it sends another message. This permits a MESSAGE PARITY ERROR message from the initiator to be associated with the appropriate message.
SELECTION	Immediately after that SELECTION phase and before the initiator releases BSY
RESELECTION	After the drive has sent its IDENTIFY message for that RESELECTION phase. The initiator should only assert the ATN signal during a RESELECTION phase to transmit a TARGET RESET or DISCONNECT message.

The initiator keeps the ATN signal asserted if more than one byte is to be transferred. The initiator can negate the ATN signal at any time, except it does not negate the ATN signal while the ACK signal is asserted during a MESSAGE OUT phase. Normally, the initiator negates the ATN signal while the REQ signal is true and the ACK signal is false during the last REQ/ACK handshake of the MESSAGE OUT phase.

B.6.2 Reset Condition

The drive employs the hard reset alternative for bus RESET processing. Write cache data is preserved across a SCSI reset and will be flushed to the media.

The hard reset means that upon detection of the reset condition, the disk drive follows these steps:

1. Clears all tasks, including task sets (queues).
2. Releases all SCSI device reservations.
3. Returns any SCSI device operating modes to their appropriate initial conditions, similar to those conditions that would be found after a normal power-on reset. MODE SELECT conditions are restored to their last saved values, if saved values have been established. MODE SELECT conditions for which no values have been saved are returned to their default values.
4. Sets unit attention condition.

B.7 SCSI MESSAGE FORMAT

A message can be one or more bytes in length. One or more messages can be sent during a single MESSAGE phase, but a message cannot be split over MESSAGE phases. The initiator is required to end the MESSAGE OUT phase (by negating ATN) when it sends certain messages.

A message cannot be split. The drive will not split messages longer than one byte even if ATN is deasserted.

When a connection to the drive is established (i.e., the drive is selected with ATN asserted), the first message byte passed by the initiator must be either an IDENTIFY, ABORT TASK SET, or TARGET RESET message. If not, the drive discards the message, saves no status information, and goes to the BUS FREE phase.

If an initiator supplies an unsupported or incorrect messages (for example, TASK COMPLETE, or a reserved or undefined message code), the drive returns a MESSAGE REJECT message and continues where it left off (possibly returning to MESSAGE OUT if ATN is raised).

The first byte of the message, as defined in Table B-10, gives the format of the message.

Table B-10 SCSI Message Format

Message Code	Message
00h	One-byte message (TASK COMPLETE)
01h	Extended message
02h – 1Fh	One-byte message
30h – 7Fh	Reserved
80h – FFh	One-byte message (IDENTIFY)

The drive supports the messages listed in Table B-11. The message code and the direction of the message flow is also included in the table.

Table B-11 Supported SCSI Messages

Message	Message Code	Direction ¹
ABORT TASK ² (ABORT TAG)	0Dh	Out
ABORT TASK SET ² (ABORT)	06h	Out
CLEAR TASK SET ² (CLEAR QUEUE)	0Eh	Out
CONTINUE TASK ² (CONTINUE I/O PROCESS)	12h	Out
DISCONNECT	04h	In Out
IDENTIFY	80h – FFh	In Out
IGNORE WIDE RESIDUE ³	23h	In
INITIATE RECOVERY	0Fh	In
INITIATOR DETECTED ERROR	05h	Out
LINKED COMMAND COMPLETE	0Ah	In
LINKED COMMAND COMPLETE (with flag)	0Bh	In
MESSAGE PARITY ERROR	09h	Out
MESSAGE REJECT	07h	In Out
NO OPERATION	08h	Out
<u>Queue Tag Messages (two bytes)</u> ³		
HEAD OF QUEUE TAG	21h	Out
ORDERED QUEUE TAG	22h	Out
SIMPLE ² (SIMPLE QUEUE TAG)	20h	In Out
RELEASE RECOVERY	10h	Out
SAVE DATA POINTER	02h	In
SYNCHRONOUS DATA TRANSFER REQUEST ⁴	01h	In Out
TARGET RESET ² (BUS DEVICE RESET)	0Ch	Out
TARGET TRANSFER DISABLE	13h	Out
TASK COMPLETE ² (COMMAND COMPLETE)	00h	In
WIDE DATA TRANSFER REQUEST ²	03h	In Out
PARALLEL PROTOCOL REQUEST ⁴	04h	In Out

Notes: ¹ In = target to initiator; Out = initiator to target (drive).

² SCSI-3 terminology (SCSI-2 term in parentheses).

³ Multibyte message.

⁴ Extended messages (see Figure B-1).

Two-byte messages consist of two consecutive bytes. The value of the first byte, as defined in Table B-10, determines which message is to be transmitted. The second byte is a parameter byte that is used as defined in the message description.

A value of 01h in the first byte indicates the beginning of a multiple-byte extended message. The extended message format is shown in Figure B-1 and the data fields are described in Table B-12.

Bit Byte	7	6	5	4	3	2	1	0
0	Extended Message (01h)							
1	Extended Message length							
2	Extended Message code							
3 to n-1	Extended Message Arguments							

Figure B-1 Extended Message — Data Format

Table B-12 Extended Message — Field Description

Field	Description
Extended Message Length	<p>This field specifies the length, in bytes, of the Extended Message Code plus the Extended Message Arguments that follow. Therefore, the total length of the message is equal to the Extended Message Length plus 2.</p> <p>A value of 0 for the Extended Message Length indicates that 256 bytes follow.</p>
Extended Message Code	<p>The drive supports three Extended Messages. They are:</p> <p>01h SYNCHRONOUS DATA TRANSFER REQUEST</p> <p>03h WIDE DATA TRANSFER REQUEST</p> <p>04h PARALLEL PROTOCOL REQUEST</p>

B.8 SUPPORTED SCSI MESSAGES

Following are descriptions of each of the messages supported by the drive. SCSI-3 message names are used. If there is a corresponding SCSI-2 name that is different, it is shown in parentheses in *ITALICS*.

B.8.1 ABORT TASK Message (0Dh) (*ABORT TAG*)

The drive's response to an ABORT TASK message depends on the task being performed when the message arrives. If anything less than an I_T_L_Q nexus is established, the drive goes to the BUS FREE phase without affecting any tasks or recording any sense data.

If an I_T_L_Q nexus exists, the drive aborts the current process and goes to the BUS FREE phase. If the drive does not have a queued or active task for the specified tag, it goes to the BUS FREE phase. Any pending data, sense data, or status for the task are cleared. However, neither an Auto Contingent Allegiance condition for the initiator nor any other tasks are affected and all cached writes are retained. Mode parameters and reservations established by earlier commands are not affected.

B.8.2 ABORT TASK SET Message (06h) (*ABORT*)

The ABORT TASK SET message is sent by the initiator to clear any task for the I_T_x nexus. The drive's response depends upon the task being performed when the message arrives.

If only an I_T nexus has been established (that is, a valid IDENTIFY message has not been received), the drive goes to the BUS FREE phase without affecting any task or recording any sense data. An exception is if the drive has received a command during this I_T nexus; the task is aborted as described below.

If an I_T_L nexus exists, the drive aborts all tasks, and goes to the BUS FREE phase. Any pending data, sense data, or status for the tasks are cleared. However, an Auto Contingent Allegiance condition for the initiator is not affected and any cached writes are retained. Mode parameters and reservations established by earlier commands are not affected.

B.8.3 CLEAR ACA Message (16h)

Since NACA bit = 1 is not supported, SCSI-2 rules for clearing ACA are in effect and this message is ignored by the drive.

B.8.4 CLEAR TASK SET Message (0Eh)

The CLEAR TASK SET message clears all tasks from all initiators and goes to the BUS FREE phase. Pending data and status for all processes are discarded. A Unit Attention condition is generated for all other initiators that had their queue commands cleared (additional sense code is set to COMMANDS CLEARED BY ANOTHER INITIATOR).

Previously established conditions, including mode parameters, reservations, and Auto Contingent Allegiance conditions are not affected. An I_T_L nexus must exist for the drive to accept this message. After accepting this message, the drive goes to the BUS FREE phase.

B.8.5 CONTINUE TASK Message (12h) (*CONTINUE I/O PROCESS*)

The CONTINUE TASK message tells the drive to resume the I/O that the initiator has reconnected. This message is sent in the same MESSAGE OUT phase as the IDENTIFY message.

B.8.6 DISCONNECT Message (04h)

The DISCONNECT message is sent from the drive to inform the initiator that the present connection is going to be broken, i.e., the drive plans to disconnect by releasing the BSY signal. A later reconnection will be required to complete the current task. The message does not cause the initiator to save the current data pointer. After sending the message, the drive goes to the BUS FREE phase by releasing BSY.

The DISCONNECT message can also be sent by the initiator to tell the drive to suspend the current phase and disconnect from the bus. The drive's response and handling of a DISCONNECT message is based on when, in the task, the initiator sends the DISCONNECT message. Table B-13 summarizes the drive's response.

Table B-13 Drive's Response to DISCONNECT Message

BUS Phase	Drive's Response
SELECTION	The drive discards the DISCONNECT message and goes to BUS FREE.
COMMAND	The ATTENTION request is ignored while the Command Descriptor Block is fetched; that is, the drive does not switch to MESSAGE OUT until after all Command Descriptor Block bytes have been received. The drive returns the DISCONNECT message, goes to BUS FREE, and reconnects to resume the task based on the disconnect/reconnect parameters.
DATA	The ATTENTION request is ignored until it is convenient to handle the condition. For example, on reads and writes, the transfer of the current host block is completed before handling the ATTENTION. The drive honors the DISCONNECT message (sends SAVE DATA POINTER and DISCONNECT messages) regardless of the amount of data remaining to be sent for the current command.
STATUS IN	The drive sends a MESSAGE REJECT message, switches back to MESSAGE IN, and sends the TASK COMPLETE message.

B.8.7 IDENTIFY Message (80h - FFh)

The IDENTIFY message is sent by either the initiator or the drive to establish an I_T_L nexus with the drive, under the conditions listed below. Figure B-2 shows the format of the IDENTIFY message and Table B-14 describes the data field contents.

Bit	7	6	5	4	3	2	1	0
	Identify	DiscPriv	LUN					

Figure B-2 IDENTIFY Message — Data Format

Table B-14 IDENTIFY Message — Field Description

Field	Description
Identify	The Identify bit must be set to 1 . This identifies this message as an IDENTIFY message.
DiscPriv	Disconnect Privilege. This bit must be set to 1 for all tagged tasks (that is, a QUEUE TAG and QUEUE TAG message follow the IDENTIFY message). If not set to 1, the drive returns BUSY status.
LUN	Only one Logical Unit Number shall be identified per task. The drive supports a single Logical Unit Number (LUN 0). A non-zero value in this field causes the drive to terminate the command with a CHECK CONDITION status, a Sense Key of ILLEGAL REQUEST, and an Additional Sense of LOGICAL UNIT NOT SUPPORTED.

B.8.8 IGNORE WIDE RESIDUE Message (23h)

The drive sends the IGNORE WIDE RESIDUE Message to indicate that the number of valid bytes sent during the last REQ/ACK handshake is less than the negotiated transfer width. The Ignore Field contains the number of invalid data bytes transferred.

B.8.9 INITIATE RECOVERY Message (0Fh)

The drive informs the initiator that it is entering an Auto Contingent Allegiance condition (ACA) by sending an INITIATE RECOVERY message immediately following either a CHECK CONDITION or a COMMAND TERMINATED status. The ACA remains in effect until it is terminated.

A MESSAGE REJECT response to an INITIATE RECOVERY message indicates that an ACA shall not be established. The enabled or disabled status of an ACA is not changed by the rejection of an INITIATE RECOVERY message.

B.8.10 INITIATOR DETECTED ERROR Message (05h)

The INITIATOR DETECTED ERROR message is sent from an initiator to inform the drive that an error has occurred that does not preclude the drive from retrying the operation. The source of the error can be related to previous activities on the SCSI bus or can be internal to the initiator and unrelated to any previous SCSI bus activity. As shown in Table B-15, the drive's response to and its handling of the message are based on when, in the task, the initiator introduces the message.

Table B-15 Drive's Response to INITIATOR DETECTED ERROR Message

BUS Phase	Drive's Response
SELECTION	The drive discards the INITIATOR DETECTED ERROR message and goes to the BUS FREE phase.
COMMAND	The drive discards any Command Descriptor Block bytes fetched from the initiator, sets the Sense Key to ABORTED COMMAND, sets the Additional Sense Code to INITIATOR DETECTED ERROR MESSAGE RECEIVED. It sends the CHECK CONDITION status and the TASK COMPLETE message and then goes to the BUS FREE phase.
DATA	The drive sets the Sense Key to ABORTED COMMAND and sets the Additional Sense Code to INITIATOR DETECTED ERROR MESSAGE RECEIVED. It sends the CHECK CONDITION status and the COMMAND COMPLETE message and then goes to the BUS FREE phase.
STATUS	The drive switches to STATUS phase and attempts to re-send the status byte to the initiator. If another INITIATOR DETECTED ERROR message is received for the same nexus while the bus is in the status phase, the drive sets the sense key to ABORTED COMMAND, sets the Additional Sense Code to INITIATOR DETECTED ERROR MESSAGE RECEIVED, and then goes to the BUS FREE phase.
MESSAGE IN	The drive switches to the MESSAGE IN phase and attempts to re-send the message to the initiator. If another INITIATOR ERROR DETECTED message is received for the same nexus while the bus is in the MESSAGE IN phase, the drive behavior depends on the message. If the drive was attempting to send the TASK COMPLETE message, it goes to the BUS FREE phase. If the drive was attempting to send any other message, it sends the CHECK CONDITION status and the TASK COMPLETE message, and then goes to the BUS FREE phase. In either case, the drive sets the Sense Key to ABORTED COMMAND and sets the Additional Sense Code to INITIATOR DETECTED ERROR MESSAGE RECEIVED.

B.8.11 LINKED COMMAND COMPLETE Message (0Ah)

The drive sends the LINKED COMMAND COMPLETE Message to the initiator to signal that the execution of a linked command has been completed and the status has been sent. The initiator should then set the pointers to the initial state for the next linked command.

B.8.12 LINKED COMMAND COMPLETE (WITH FLAG) Message (0Bh)

The drive sends the LINKED COMMAND COMPLETE (WITH FLAG) message to the initiator to signal that the execution of a linked command (with the flag bit set) has been completed and the status has been sent. The initiator should then set the pointers to the initial state for the next linked command.

B.8.13 MESSAGE PARITY ERROR Message (09h)

This message tells the drive that the last message byte it passed to the host contained a parity error. If the bus phase was not Message In when the host asserted ATN, the drive releases the SCSI bus.

B.8.14 MESSAGE REJECT Message (07h)

The MESSAGE REJECT message is sent from either the drive or the initiator to indicate that the last message or message byte received was inappropriate or has not been implemented.

To indicate its intention of sending this message, the initiator asserts the ATN signal before it releases the ACK signal for the REQ/ACK handshake of the message byte that is to be rejected. If the drive receives this message under any other circumstances, it rejects the message

When a drive sends the MESSAGE REJECT message, it changes to the MESSAGE IN phase and sends it before requesting additional messages bytes from the initiator. This provides an interlock so that the initiator can determine which message byte was rejected.

After a drive sends a MESSAGE REJECT message, and if the ATN signal is still asserted, the drive returns to the MESSAGE OUT phase. The subsequent MESSAGE OUT phase begins with the first byte of a message.

B.8.15 NO OPERATION Message (08h)

The NO OPERATION message allows an initiator to tell the drive to continue with the operation it was performing before the drive communicated a need for a MESSAGE OUT phase.

For example, if the drive does not respond to an ATN condition until a later phase, and at that time the original message is no longer valid, the initiator can send the NO OPERATION message when the drive enters the MESSAGE OUT phase.

B.8.16 Queue Tag Messages (20h, 21h or 22h)

Queue Tag messages are a part of SCSI command queuing. They establish the I_T_L_Q nexus with the drive. If an IDENTIFY message does not immediately precede the Queue Tag message, the drive discards the message and goes to the BUS FREE phase.

If a valid IDENTIFY message precedes the Queue Tag message but does not grant the disconnect privilege, the drive fetches the Command Descriptor Block and immediately returns BUSY status.

Queue Tag messages are two-byte SCSI messages. The message format is shown in Figure B-3.

Bit	7	6	5	4	3	2	1	0
Byte 0	Message Code (20h, 21h, or 22h)							
Byte 1	Queue Tag							

Figure B-3 Queue Tag Message — Data Format

Command queuing is enabled or disabled by the DQue bit in the MODE SELECT Control Mode page. If that bit has queuing disabled, the drive returns a MESSAGE REJECT message in response to a QUEUE TAG message. Since an I_T_L_Q nexus was not established, the I_T_L nexus remains and any commands received are treated as untagged.

The drive supports all three SCSI-defined queue tags:

- HEAD OF QUEUE TAG
- ORDERED QUEUE TAG
- SIMPLE (*SIMPLE QUEUE TAG*)

The type of queue tag specified affects the execution order of the commands. The tags specify the placement of a task in a task queue. The execution orders apply to commands from all initiators.

B.8.16.1 HEAD OF QUEUE TAG Message (21h)

The HEAD OF QUEUE TAG message specifies that the task is to be placed first in the drive's task queue. An in-progress task is not pre-empted by the command that accompanies this message. A later task, also received with a HEAD OF QUEUE TAG message, would be placed at the head of the drive's task queue for execution in a last-in, first-out order.

B.8.16.2 ORDERED QUEUE TAG Message (22h)

The ORDERED QUEUE TAG message specifies that the task be placed in the drive's task queue for execution in the order received. All tasks placed in the task queue before this process is received get executed before this process. Except for tasks received with a HEAD OF QUEUE TAG message, all later tasks are executed after this task.

B.8.16.3 SIMPLE Message (20h) (*SIMPLE QUEUE TAG*)

The SIMPLE message specifies that the task be placed in the drive's task queue. If command reordering is enabled, commands queued with a SIMPLE QUEUE TAG may be executed in an order that optimizes overall drive throughput, though the execution order will observe the HEAD OF QUEUE TAG and ORDERED QUEUE TAG message order rules.

B.8.17 RELEASE RECOVERY Message (10h)

The RELEASE RECOVERY message terminates an Auto Contingent Allegiance condition established by a previous INITIATE RECOVERY message. See Chapter 5 for a description of the drive's processing done upon Auto Contingent Allegiance termination. An IDENTIFY Message must immediately precede this message.

The drive rejects the message if there is no Auto Contingent Allegiance condition active when the message is received or if the host fails to precede it with an IDENTIFY message.

B.8.18 SAVE DATA POINTER Message (02h)

The SAVE DATA POINTER message is sent from the drive to direct the initiator to copy the current data pointer to the saved data pointer for the current task. The message is sent prior to the drive performing a bus disconnect.

B.8.19 SYNCHRONOUS DATA TRANSFER REQUEST Message (01h)

The SYNCHRONOUS DATA TRANSFER REQUEST (SDTR) message exchange establishes the permissible transfer period and the REQ/ACK offsets for the DATA IN/DATA OUT phases. The transfer period is the minimum time allowed between leading edges of successive REQ and ACK pulses. The REQ/ACK offset is the maximum number of REQ pulses allowed to be outstanding before the leading edge of its corresponding ACK pulse is received by the drive.

The agreement becomes invalid and reverts to asynchronous data transfer after any condition that can leave the data transfer agreement in an indeterminate state such as:

- after a hard reset,
- after a TARGET RESET message,
- after a WIDE DATA TRANSFER REQUEST message, and
- after a power cycle.

The format of the message fields are shown in Figure B-4 and described in Table b-16.

Bit Byte	7	6	5	4	3	2	1	0
0	Extended Message (01h) (See Figure B-1)							
1	Extended Message Length (03h)							
2	SYNCHRONOUS DATA TRANSFER REQUEST code (01h)							
3	Transfer Period							
4	REQ/ACK Offset							

Figure B-4 SYNCHRONOUS DATA TRANSFER-Data Format

Table B-16 SYNCHRONOUS DATA TRANSFER REQUEST Message —
Field Descriptions

Field	Description
Transfer Period	Transfer periods are specified in 4-ns units (a transfer period value of 100 yields an actual transfer period of 400 ns (4 x 100)). The minimum Transfer Period is 50 ns for ULTRA SCSI-3. The drive supports transfer periods in the range from 50 to 400 ns in increments of 25 ns (that is 50, 75, 100...400). (The equivalent transfer period value (entries) for these are 12, 18, 25...100, respectively.) The drive rounds up the initiator-supplied transfer period that does not fall on one of these intervals
REQ/ACK Offset	<p>The drive supports a minimum offset value of 0. An offset value of zero means that all data transfers are done asynchronously. The drive's maximum and default offset value is 31. If the initiator supplies an offset value of less than 31 in its SYNCHRONOUS DATA TRANSFER REQUEST message, the drive rounds down to the initiator's value and echoes it in its corresponding SYNCHRONOUS DATA TRANSFER REQUEST message; otherwise, the drive returns a value of 31 as its synchronous data transfer offset.</p> <p>The drive performs asynchronous data transfers by default (after a hard reset, the receipt of a SYNCHRONOUS DATA TRANSFER REQUEST message exchange must occur before synchronous data transfers can take place). The drive or the initiator can initiate the SYNCHRONOUS DATA TRANSFER REQUEST message exchange based on the Initiate Sync Negotiation (ISN) parameter For details, see section of this manual that discusses the MODE SELECT commands "Quantum-Unique Page" (39h).</p>

B.8.20 TARGET RESET Message (0Ch) (BUS DEVICE RESET)

The TARGET RESET message is sent from an initiator to direct the drive to clear all tasks on the drive. Refer back to B.6.2, Reset Condition, for the steps taken during this process.

The drive creates a Unit Attention condition for all initiators after accepting and processing a TARGET RESET message. The additional sense code is set to TARGET RESET OCCURRED.

B.8.21 TARGET TRANSFER DISABLE Message (13h)

The TARGET TRANSFER DISABLE message tells the drive that subsequent reconnections for data transfer on the task must be performed by the initiator. The target can connect for other purposes but cannot enter a DATA phase unless a disconnection attempt is rejected.

When the drive is ready to transfer data for a disconnected task for which a TARGET TRANSFER DISABLE message has been sent, the target reconnects to the initiator for the task, sends a DISCONNECT message, and, if the initiator does not respond with a MESSAGE REJECT message, goes to the BUS FREE phase.

B.8.22 TASK COMPLETE Message (00h) (COMMAND COMPLETE)

The TASK COMPLETE message is sent by the drive to an initiator to indicate that a task has been completed and that valid status has been sent to the initiator. When

the drive detects the negation of ACK for the TASK COMPLETE message (with ATN false), the drive goes to the BUS FREE phase by releasing the BSY signal. The drive also negates the MSG, C/D and I/O signals at this time.

B.8.23 WIDE DATA TRANSFER REQUEST Message (03h)

The WDTR message exchange establishes an agreement between two SCSI devices on the width of the data path used for DATA phase transfers between the two devices. This agreement applies to DATA IN and DATA OUT phases only. All other information transfer phases use an 8-bit data path.

If a SCSI device implements the wide data transfer option and the synchronous data transfer option, then the device negotiates the wide data transfer agreement before negotiating the synchronous data transfer agreement. If a synchronous data transfer agreement is in effect, the SCSI device that accepts the WDTR message resets the synchronous agreement to the asynchronous mode.

The agreement becomes invalid and reverts to an 8-bit data transfer after any condition that can leave the data transfer agreement in an indeterminate state such as:

- After a hard reset,
- After a TARGET RESET message, and
- After a power cycle.

The format of the WIDE DATA TRANSFER REQUEST message is shown in Figure B-5. The data is described in Table B-17.

Bit	7	6	5	4	3	2	1	0
0	Extended Message (01h) (See Figure B-1)							
1	Extended Message Length (02h)							
2	WIDE DATA TRANSFER REQUEST code (03h)							
3	Transfer Width							

Figure B-5 WIDE DATA TRANSFER REQUEST Message — Data Format

Table B-17 WIDE DATA TRANSFER REQUEST Message — Field Descriptions

Field	Description
Transfer Width	<p>The transfer width that is established applies to all logical units on both SCSI devices. The originating SCSI device sets its transfer width value to the maximum data path width that it elects to use. If the responding SCSI device can also accommodate this transfer width, it returns the same value in its WDTR message. If it requires a smaller transfer width, it substitutes the smaller value in its message. Valid Transfer Width values are:</p> <p>00h 8-bit data transfer path (standard data path width) 01h 16-bit data transfer path (Wide SCSI option)</p> <p>The drive performs 8-bit data transfers by default. After a hard reset, the receipt of a WIDE DATA TRANSFER REQUEST message exchange must take place before wide data transfers can take place. The drive or the initiator can initiate the WIDE DATA TRANSFER REQUEST message exchange based on the Initiate Wide Negotiation Mode Bit (see the section of this manual that discusses the MODE SELECT command's "Quantum-Unique Page," [39h]).</p>

B.8.24 PARALLEL PROTOCOL REQUEST Message (04h)

The PARALLEL PROTOCOL REQUEST messages are used to communicate a synchronous data transfer agreement, a wide data transfer agreement, and set the protocol options between two SCSI devices (Figure B-6).

Bit	7	6	5	4	3	2	1	0
0	Extended Message (01h)							
1	Extended Message Length (06h)							
2	Parallel Protocol Request (04h)							
3	Transfer Period Factor							
4	Reserved							
5	REQ / ACK Offset							
6	Transfer Width Exponent (m)							
7	Reserved					QAS_Req	DT_Req	IU_Req

NOTE: Bits 0 – 3 of byte 7 (IU_Req, DT_Req, and QAS_Req) are collectively known as the Protocol Option bits, and are defined as such in Table B-18.

Figure B-6 PARALLEL PROTOCOL REQUEST Message — Data Format

B.8.24 PARALLEL PROTOCOL REQUEST Message (04h) (continued)

Table B-18 PARALLEL PROTOCOL REQUEST Message — Field Descriptions

Field	Description																		
Transfer Period Factor	<p>Options are:</p> <table border="1"> <thead> <tr> <th>Code</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>00h – 08h</td> <td>Reserved for faster timings that may be allowed in future SCSI interface standards.</td> </tr> <tr> <td>09h</td> <td>Transfer period = 12.5ns (Fast-80 data is latched every 12.5ns); valid only if the PROTOCOL OPTIONS field has a value selected that supports double-transition data transfers.</td> </tr> <tr> <td>0Ah</td> <td>Transfer period = 25ns (Fast-40 data is latched every 25ns)</td> </tr> <tr> <td>0Bh</td> <td>Transfer period = 30.3ns (Fast-40 data is latched every 30.3ns)</td> </tr> <tr> <td>0Ch</td> <td>Transfer period = 50ns (Fast-20 data is latched using a transfer period of less than or equal to 96ns and greater than or equal to 50ns)</td> </tr> <tr> <td>0Dh – 18h</td> <td>Transfer period = the period factor x 4 (Fast-20 data is latched using a transfer period of less than or equal to 96ns and greater than or equal to 50ns)</td> </tr> <tr> <td>19h – 31h</td> <td>Transfer period = the period factor x 4 (Fast-10 data is latched using a transfer period of less than or equal to 196ns and greater than or equal to 100ns)</td> </tr> <tr> <td>32h – FFh</td> <td>Transfer period = the period factor x 4 (Fast-5 data is latched using a transfer period of less than or equal to 1020ns and greater than or equal to 200ns)</td> </tr> </tbody> </table> <p>For best performance for an Atlas 10K Ildisk drive, the value of this field is 09h</p>	Code	Description	00h – 08h	Reserved for faster timings that may be allowed in future SCSI interface standards.	09h	Transfer period = 12.5ns (Fast-80 data is latched every 12.5ns); valid only if the PROTOCOL OPTIONS field has a value selected that supports double-transition data transfers.	0Ah	Transfer period = 25ns (Fast-40 data is latched every 25ns)	0Bh	Transfer period = 30.3ns (Fast-40 data is latched every 30.3ns)	0Ch	Transfer period = 50ns (Fast-20 data is latched using a transfer period of less than or equal to 96ns and greater than or equal to 50ns)	0Dh – 18h	Transfer period = the period factor x 4 (Fast-20 data is latched using a transfer period of less than or equal to 96ns and greater than or equal to 50ns)	19h – 31h	Transfer period = the period factor x 4 (Fast-10 data is latched using a transfer period of less than or equal to 196ns and greater than or equal to 100ns)	32h – FFh	Transfer period = the period factor x 4 (Fast-5 data is latched using a transfer period of less than or equal to 1020ns and greater than or equal to 200ns)
Code	Description																		
00h – 08h	Reserved for faster timings that may be allowed in future SCSI interface standards.																		
09h	Transfer period = 12.5ns (Fast-80 data is latched every 12.5ns); valid only if the PROTOCOL OPTIONS field has a value selected that supports double-transition data transfers.																		
0Ah	Transfer period = 25ns (Fast-40 data is latched every 25ns)																		
0Bh	Transfer period = 30.3ns (Fast-40 data is latched every 30.3ns)																		
0Ch	Transfer period = 50ns (Fast-20 data is latched using a transfer period of less than or equal to 96ns and greater than or equal to 50ns)																		
0Dh – 18h	Transfer period = the period factor x 4 (Fast-20 data is latched using a transfer period of less than or equal to 96ns and greater than or equal to 50ns)																		
19h – 31h	Transfer period = the period factor x 4 (Fast-10 data is latched using a transfer period of less than or equal to 196ns and greater than or equal to 100ns)																		
32h – FFh	Transfer period = the period factor x 4 (Fast-5 data is latched using a transfer period of less than or equal to 1020ns and greater than or equal to 200ns)																		
REQ / ACK Offset	<p>Request / Acknowledge Offset. For single transition (ST) synchronous data transfers, the value of REQ/ACK Offset is the maximum number of REQ assertions allowed to be outstanding before a related ACK assertion is received at the target. The size of the data transfer may be 1 or 2 bytes, depending on the values in the Transfer Width Exponent Field.</p> <p>For dual transition (DT) synchronous data transfers, the value of REQ/ACK Offset is the maximum number of REQ assertions allowed to be outstanding before a related ACK assertion is received at the target. The size of the data transfer will be 2 bytes.</p> <p>For best performance for an Atlas 10K Ildisk drive, the value of this field is 1Fh.</p>																		

(continued)

B.8.24 PARALLEL PROTOCOL REQUEST Message (04h) (continued)

Table B-18 PARALLEL PROTOCOL REQUEST Message — Field Descriptions (continued)

Field	Description												
Transfer Width Exponent	<p>The value within this field defines the width of the transfer used during data-in buffer transfers and data-out buffer transfers. The transfer width that is established applies to all logical units on both SCSI devices. Valid transfer widths are 8 bits (m = 00h) and 16 bits (m = 01h), if the Protocol Options field value is 0h; for any other value in the Protocol Options field, the only valid transfer width is 16 bits.</p> <p>For best performance for an Atlas 10K IIdisk drive, the value of this field is 01h.</p>												
Protocol Option Bits (IU_REQ, DT_REQ, and QAS_REQ)	<p>This bits are used by the originating SCSI device to indicate which protocol options are to be enabled. The SCSI device responding uses these bits to indicate which of the requested options have been enabled.</p> <p>An Information Units Enable Request (IU_REQ) = 0 indicates that information units are to be disabled when received from the originator and that information units received from the responding SCSI device are not supported. An IU_REQ bit = 1 indicates that information units are to be enabled when received from the originator and that information units received from the responding device are supported.</p> <p>A Dual Transition Request (DT_REQ) = 0 indicates DT DATA phases are to be disabled when received from the originator and that information units received from the responding SCSI device are not supported. An DT_REQ bit = 1 indicates that DT DATA phases are to be enabled when received from the originator and that DT DATA phases received from the responding device are supported.</p> <p>A QAS Enable Request (QAS_REQ) = 0 indicates that QAS is to be disabled when received from the originator and QAS received from the responding SCSI device is not supported. An QAS_REQ bit = 1 indicates that QAS is to be enabled when received from the originator and that QAS received from the responding device are supported.</p> <p>The Atlas 10K II disk drive supports two operating modes selected via the Protocol Option Bits:</p> <table border="1"> <thead> <tr> <th>QAS_REQ</th> <th>DT_REQ</th> <th>IU_REQ</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>0</td> <td>0</td> <td>Use ST data-in buffer transfers and ST data-out buffer transfers to transfer data.</td> </tr> <tr> <td>0</td> <td>1</td> <td>0</td> <td>Use DT data-in buffer transfers and DT data-out buffer transfers to transfer data.</td> </tr> </tbody> </table> <p>For best performance for an Atlas 10K II disk drive the value of Byte 7 should be 02h.</p>	QAS_REQ	DT_REQ	IU_REQ	Description	0	0	0	Use ST data-in buffer transfers and ST data-out buffer transfers to transfer data.	0	1	0	Use DT data-in buffer transfers and DT data-out buffer transfers to transfer data.
QAS_REQ	DT_REQ	IU_REQ	Description										
0	0	0	Use ST data-in buffer transfers and ST data-out buffer transfers to transfer data.										
0	1	0	Use DT data-in buffer transfers and DT data-out buffer transfers to transfer data.										

NEGOTIATED RATE INFORMATION PAGE REFERENCE

This appendix provides information about the INQUIRY command's Negotiated Rate Information Page that is not provided in Section 5.11.2.7. Specifically, it provides the codes and meanings for the Transfer Period Factor field (byte 4) and a table that explains the allowable settings for the QAS, DT, and IU fields (bits 2, 1, and 0, respectively, of byte 7).

C.1 TRANSFER PERIOD FACTOR FIELD

Table C-1 provides the codes and meanings allowed for the Transfer Period Factor field if the value of PARL (Parallel Protocol, bit 2 of byte 6 of the Negotiated Rate Information Page) is 1.

Table C-2 provides the allowable codes when the value of PARL = 0.

Table C-1 Transfer Period Factor Field Values When PARL = 1

Code	Description
00h – 08h	Reserved.
09h	Transfer period = 12.5 ns (Fast-80 is latched every 12.5 ns). This code is valid only if the protocol options bits (QAS, DT, and IU) of the page have values selected that support double transition (DT) data transfers.
0Ah	Transfer period = 25 ns (Fast-40 data is latched every 25 ns).
0Bh	Transfer period = 30 ns (Fast-40 data is latched every 30 ns).
0Ch	Transfer period = 50 ns (Fast-20 data is latched using a transfer period of less than or equal to 96 ns and greater than or equal to 50 ns).
0Dh – 18h	Transfer period = the period factor x 4 (Fast-20 data is latched using a transfer period of less than or equal to 96 ns and greater than or equal to 50 ns).
19h – 31h	Transfer period = the period factor x 4 (Fast-10 data is latched using a transfer period of less than or equal to 196 ns and greater than or equal to 100 ns).
32h - FFh	Transfer period = the period factor x 4 (Fast-5 data is latched using a transfer period of less than or equal to 1020 ns and greater than or equal to 200 ns).

Table C-2 Transfer Period Factor Field Values When PARL = 0

Code	Description
00h – 09h	Reserved.
0Ah	Transfer period = 25 ns (Fast-40 data is latched every 25 ns).
0Bh	Transfer period = 30 ns (Fast-40 data is latched every 30 ns).
0Ch	Transfer period = 50 ns (Fast-20 data is latched using a transfer period of less than or equal to 96 ns and greater than or equal to 50 ns).
0Dh – 18h	Transfer period = the period factor x 4 (Fast-20 data is latched using a transfer period of less than or equal to 96 ns and greater than or equal to 50 ns).
19h – 31h	Transfer period = the period factor x 4 (Fast-10 data is latched using a transfer period of less than or equal to 196 ns and greater than or equal to 100 ns).
32h - FFh	Transfer period = the period factor x 4 (Fast-5 data is latched using a transfer period of less than or equal to 1020 ns and greater than or equal to 200 ns).

C.2 SETTINGS FOR THE QAS, DT, AND IU FIELDS

Table C-3 provides valid bit combinations for the QAS, DT, and IU fields of the INQUIRY command's Negotiated Rate Information Page (Section 5.11.2.7 of this manual). Only the combinations shown in the table are valid; all other combinations are reserved.

Table C-3 Settings for the QAS, DT, and IU Fields of the Negotiated Rate Information Page (INQUIRY Command)

QAS	DT	IU	Description
0	0	0	Selects Single Transition phases. Use ST DATA IN and ST DATA OUT phases to transfer data
0	1	0	Use DT DATA IN and DT DATA OUT phases with data group transfers.
0	1	1	Use DT DATA IN and DT DATA OUT phases with information unit transfers.
1	1	1	Use DT DATA IN and DT DATA OUT phases with information unit transfers and use QAS for arbitration.

INDEX

—A—

ABORT TASK Message, B-28
ABORT TASK SET Message, B-28
Acoustic Emissions, 4-10
Activity LED, 3-14
Additional Sense Codes, 5-179
Additional Sense Length, 5-173, 5-174
AIP, starting, 5-98
ANSI Version, 5-35, 5-36
ARBITRATION Phase, B-13
ASCII Implemented Operating Definition Page,
5-37
Asynchronous Data Transfer, B-18
Attention Condition, B-23

—B—

Bit Pointer, 5-187
Block Format Defect Descriptor, 5-26
Block Length in Bytes, 5-141, 5-142
Buffer Empty Ratio, 5-68, 5-69
Buffer Full Ratio, 5-68, 5-69
Buffer ID, 5-139, 5-140, 5-221
Buffer Offset, 5-139, 5-140, 5-221
Build Date and Time, 5-43
Bus Inactivity Limit, 5-68, 5-69, 5-70
Bus Phases, B-12

—C—

Cables, 3-28
Cache Table Size, 5-110
Caching Page, 5-74
CHANGE DEFINITION Command, 5-23
Changeable Pages, 5-60, 5-61, 5-63
CLEAR ACA Message, B-29
CLEAR QUEUE, 5-9
CLEAR TASK SET Message, B-29
Clearing Sense Data, 5-172
COMMAND COMPLETE Message, 5-13
Configuration Jumpers and Connections, 3-6
Configuring SCSI ID (Ultra 160/m SCSI
Versions with 68-Pin Connector), 3-8
Configuring SCSI ID for SCA-2 Connector
Versions, 3-12
Connect Time Limit, 5-68, 5-69
CONNECTOR PINOUTS AND WIRING DIAGRAMS,
3-17
CONTINUE TASK Message, B-29
Control Mode Page, 5-76
Correction Span, 5-65, 5-67, 5-74
Cylinder Number of Defect, 5-30, 5-146

Cylinder Skew Factor, 5-107, 5-108

—D—

Data Strobe Offset Count, 5-65
Data Verification, 6-5
Device Identification Page, 5-37, 5-41
Device Type Modifier, 5-35
DISCONNECT message, B-30
Disconnect Time Limit, 5-68, 5-69
Disconnect-Reconnect Mode Page, 5-60, 5-68
Disk Controller Hardware Revision Number,
5-35
Drive Connections, 3-27
Drive Installation, 3-1
Drive Serial Number, 5-35, 5-41
Drive Step Rate, 5-107, 5-108

—E—

ECMA Version, 5-35, 5-47
Electronics Pass Number, 5-35
Enable Narrow Mode (Force 8), 3-11
Ending Boundary, 5-79, 5-81
Error Reporting, 5-9
Extent List Length, 5-191, 5-192

—F—

Field Pointer Bytes, 5-188, 5-189
Firmware Revision Page, 5-37, 5-42
Force 8, 3-11
Force Single-Ended Operation, 3-15
Format Device, 5-104
Format Device Page, 5-105
FORMAT UNIT Command, 5-25

—G—

GList, 5-141, 5-143
Grown Defect List, 5-25, 5-26, 5-27, 5-28, 5-66, 5-143

—H—

HARDWARE ERROR, 5-175, 5-177
Head Number of Defect, 5-30, 5-144
HEAD OF QUEUE TAG Message, B-35
Hot Plugging, 6-11

—I—

IDENTIFY Message, B-30
IGNORE WIDE RESIDUE Message, B-31
Immed, 5-28

Implemented Operating Definition Page, 5-36, 5-38
Information Transfer Phases, B-17
Initialization Pattern Descriptor, 5-27, 5-28, 5-30
INITIATE RECOVERY Message, B-31
INITIATOR DETECTED ERROR Message, B-31
INQUIRY Command, 5-33
Interleave, 5-26, 5-105, 5-106
ISO Version, 5-35, 5-47

—J—

Jumper Configurations and Connections – Ultra 160/m
SCSI drives with 68-Pin SCSI Connector),
3-6
Jumper Configurations and Connections for SCA-2
Connector Versions, 3-12

—L—

Landing Zone Cylinder, 5-107, 5-108
LINKED COMMAND COMPLETE (WITH FLAG) Message,
B-32
LINKED COMMAND COMPLETE Message, B-32
LOG SELECT Command, 5-51
LOG SENSE Command, 5-53
LVD DIFFSENS Driver, 6-10
LVD DIFFSENS Receiver, 6-10
LVD/SE Multimode and HVD, 6-10

—M—

Maximum Burst Size, 5-68, 5-70
Maximum Number of Notches, 5-79, 5-80
Media Error Protection, 6-4
Medium Rotation Rate, 5-107, 5-109
Medium Type, 5-62, 5-99, 5-104, 5-111
MESSAGE PARITY ERROR Message, B-33
MESSAGE REJECT Message, B-33
MLVD, 6-9
Mode Data Length, 5-62, 5-92, 5-104, 5-111
Mode Page Types, 5-61
Mode Parameter Block Descriptor, 5-58
Mode Parameter Header, 5-61, 5-62, 5-101
Mode Parameter List, 5-61
MODE SELECT (6) Command, 5-59
MODE SELECT (10) Command, 5-99
MODE SENSE (6) Command, 5-103
MODE SENSE (10) Command, 5-109
Mounting, 3-21, 3-22, 3-24

—N—

Negotiated Rate Information Page, 5-45
New Operating Definition, 5-23
NO OPERATION Message, B-34
NO SENSE, 5-92, 5-171
Non-operating Environment, 4-11
Normal commands, 5-3

NOT READY, 5-86, 5-175, 5-180, 5-181
Notch and Partition Page, 5-60, 5-79
Number of Cylinders, 5-107, 5-108
Number of Heads, 5-107, 5-108

—O—

ORDERED QUEUE TAG Message, B-35

—P—

Pages Notched, 5-81
PARALLEL PROTOCOL REQUEST Message, B-40
Parameter Code Reset, 5-51
Parameter Pointer, 5-55
Parameter Pointer Control, 5-55
Periodic Self-Adjustments, 6-7
PERSISTENT RESERVATION IN Command, 5-113
PERSISTENT RESERVATION OUT Command, 5-123
PList, 5-143, 5-145
Power Condition Page, 5-60, 5-84
Power Management, 6-6
Power Requirements, 4-7
Power-Up Self-Test, 6-7
Prefetch, 6-2
Primary Defect List, 5-25, 5-27, 5-143, 5-145, 5-147
Progress Indication, 5-171, 5-188, 5-189, 5-190
Progress Indication Bytes, 5-171, 5-189, 5-190

—Q—

Quantum Manufacturing Information Page, 5-43
Quantum-Unique Page, 5-60, 5-91
Queue Algorithm Modifier, 5-76, 5-77
Queue Tag Messages, B-34

—R—

Reallocation of Bad Blocks, 6-8
READ (6) Command, 5-135
READ (10) Command, 5-137
READ BUFFER Command, 5-139
READ CAPACITY Command, 5-141
READ DEFECT DATA Command, 5-143
READ LONG Command, 5-151
Read Retry Count, 5-65, 5-67
Read-Write Error Recovery Page, 5-65
Ready AEN Holdoff Period, 5-76
READ SKIP MASK Command, 5-153
REASSIGN BLOCKS Command, 5-155
RELEASE (6) Command, 5-163
RELEASE (10) Command, 5-165
RELEASE RECOVERY Message, B-35
Remote Busy and Fault Displays, 3-11
REPORT DEVICE IDENTIFIER Command, 5-167
REPORT LUNS Command, 5-169
REQUEST SENSE Command, 5-171
RESELECTION Phase, B-15

- Reservation Identification, 5-163, 5-165, 5-166, 5-191, 5-192, 5-193
- RESERVE (6) Command, 5-191
- RESERVE (10) Command, 5-193 —U—
- Reset Condition, B-24
- Response Data Format, 5-35, 5-36
- Retry Count, 5-188, 5-190
- Returned Logical Block Address, 5-141, 5-142 —V—
- REZERO UNIT Command, 5-197
- Rigid Disk Geometry Page, 5-107
- Rotational Offset, 5-107, 5-109
- S—
- SAVE DATA POINTER Message, B-35
- SCSI Bus Conditions, B-23
- SCSI Hardware Revision Number, 5-35
- SCSI Message Format, B-25
- SEEK (6) Command, 5-199
- SEEK (10) Command, 5-201
- SELECTION Phase, B-14
- SEND DIAGNOSTIC Command, 5-203
- Sense Data Availability, 5-171
- Sense Keys, 5-175
- SET DEVICE IDENTIFIER Command, 5-205
- Signal Restrictions Between Phases, B-21
- SIMPLE Message, B-35
- SPECIFICATIONS, 4-1
- Spin Up (Ultra 160/m SCSI drives with 68-Pin Connector), 3-10
- Spin Up for SCA-2 Connector Versions, 3-13
- Standard Inquiry Data Page, 5-35
- Starting Boundary, 5-79, 5-81
- START/STOP UNIT Command, 5-207
- Status Reporting, 5-172
- Supplied Format, 5-161, 5-162, 2-204
- Supported Additional Sense Codes, 5-179
- Supported Diagnostic Page List, 5-204
- Supported Page List, 5-37
- Supported Sense Keys, 5-175
- Supported Vital Product Data Pages Page, 5-37
- SYNCHRONIZE CACHE Command, 5-209
- SYNCHRONOUS DATA TRANSFER REQUEST Message, B-36
- T—
- TARGET RESET Message, B-37
- TARGET TRANSFER DISABLE Message, B-37
- TASK COMPLETE, B-38
- TERMPWR (Ultra 160/m SCSI drives with 68-Pin Connector), 3-8
- TEST UNIT READY Command, 5-211
- Track Skew Factor, 5-105, 5-106
- Translate Format, 5-161, 5-162, 5-204
- Translated Address (Cylinder Number), 5-161
- Translated Address (Head Number), 5-161
- Unit Attention Condition Page, 5-64
- V—
- Vendor Specific, 5-136, 5-142, 5-226
- Vendor Specific Bits, 5-8
- Vendor-Unique Caching Page, 5-110
- Verification Length, 5-209
- VERIFY Command, 5-213
- Verify Correction Span, 5-71, 5-72
- Verify Recovery Time Limit, 5-71, 5-72
- Verify Retry Count, 5-71, 5-72
- Vital Product Data Pages, 5-39
- W—
- WIDE DATA TRANSFER REQUEST Message, B-38
- WRITE (6) Command, 5-215
- WRITE (10) Command, 5-217
- WRITE AND VERIFY Command, 5-219
- WRITE BUFFER Command, 5-221
- WRITE LONG Command, 5-223
- Write Protection (Ultra 160/m SCSI drives with 68-Pin Connector), 3-10
- Write Retry Count, 5-65, 5-67
- WRITE SAME Command, 5-225
- WRITE SKIP MASK Command, 5-227
- X—
- XDREAD (10) Command, 5-229
- XDWRITE (10) Command, 5-231
- XOR Control Mode Page, 5-60, 5-82
- XPWRITE (10) Command, 5-233

Quantum:

500 McCarthy Blvd.
Milpitas, CA 95035
(408) 894-4000

www.quantum.com

June 8, 2000
81-122517-04