



.....
Disc Drive SCSI-2/SCSI-3 Interface
.....

Family Models:
.....

ST12400, ST11900, ST31200, ST11950
.....

ST12450, ST15150, ST3655, ST3390
.....

ST3285, ST31230, ST31231, ST32430
.....

ST15230, ST31250, ST32151, ST31051
.....

ST32550, ST34371, ST34501, ST32171
.....

ST32155, ST31055, ST410800, ST18771
.....

ST19171, ST12550, ST19101, ST423451
.....

ST118273, ST34501, ST52160, ST39173
.....

ST34573, ST34572, ST32272, ST39102
.....

ST118202
.....

Product Manual, Volume 2; Version 2
.....

.....
Disc Drive SCSI-2/SCSI-3 Interface
.....

Family Models:
.....

ST12400, ST11900, ST31200, ST11950
.....

ST12450, ST15150, ST3655, ST3390
.....

ST3285, ST31230, ST31231, ST32430
.....

ST15230, ST31250, ST32151, ST31051
.....

ST32550, ST34371, ST34501, ST32171
.....

ST32155, ST31055, ST410800, ST18771
.....

ST19171, ST12550, ST19101, ST423451
.....

ST118273, ST34501, ST52160, ST39173
.....

ST34573, ST34572, ST32272, ST39102
.....

ST118202
.....

Product Manual, Volume 2; Version 2
.....



© 1997 Seagate Technology, Inc. All rights reserved
Publication Number: 77738479, Rev H
August 1997

Seagate, Seagate Technology, and the Seagate logo are registered trademarks of Seagate Technology, Inc. SeaFAX, SeaFONE, SeaNET, SeaTDD, and SeaBOARD are either trademarks or registered trademarks of Seagate Technology, Inc. or one of its subsidiaries. All other trademarks or registered trademarks are the property of their respective owners.

Seagate reserves the right to change, without notice, product offerings or specifications. No part of this publication may be reproduced in any form without written permission of Seagate Technology, Inc.

Revision Status Summary Sheet

Revision	Authority	Date	Signature	Sheets Affected
A Issue		12/18/92	D. Ashby/ J. Averyt	1 thru 257.
B	PLD: 83345	8/18/93		1 thru 4, 4.1, 5 thru 193. deleted shs 194 thru 257. Technical changes on pages 5, 9, 10, 78, 93, 94, 104 - 107, 109, 136, 139 - 142, 144, 147, 149, 161 - 164, 166, 167, 181, 185 thru 193.
C	PLD: 83534	93 Oct. 19	SS	Revised sheets v thru ix, 1 thru 193. (Technical changes on pages v, vii, 1, 4, 5, 7, 8, 47, 70, 73, 78, 79, 83, 84, 86, 87, 90, 91, 93, 94, 104, 105, 107, 112, 114, 119, 123-131, 136-144, 147-156, 158-160, 162-170, 176, 178-180, 184, 185, 189, 192 and 193).
D	PLD: 85300	5/26/94	SS	v thru ix, 1 thru 189; (technical changes on pages 1, 8, 12, 26, 29, 34, 35, 70, 71, 77, 79, 80, 84, 85, 86, 95, 101, 104, 125, 150, 125, 150, 151, 160, 162, 163 & 164.)
E	PLD: 87127	01/24/95	QB	v thru ix, 1 thru 187; deleted shs 188 and 189. Renumbered pages 94 thru 187. (Technical changes on pages vi, vii, ix, 1, 2, 18, 21, 25, 26, 32, 33, 47, 50, 57, 68, 72, 74, 75, 77, 78, 79, 87, 88 - 90, 94, 105, 110, 114, 122, 124, 126, 127, 133, 135, 136, 139, 140, 143, 150, 152, 154, 158, 160, 162 - 164, 168 - 175, 178, 187.
F	PLD: 86464	06/20/95	QB	v thru ix, 1 through 187; added pages 188 through 197.
G	PLD: 86785	05/02/96	QB	v thru ix, 1 through 197; added 198 - 213 and A-1.
H	PLD: 91095	08/18/97	QB	Revised and retraced all sheets, added 214-216. (Technical changes on pages 85, 89, 97, 101, 126, 145, 152, 153, 157, 165, 168-170, 176, 177, 180, 181, 196, 198, 200-203.)

Table of contents

1.0	Interface requirements	1
1.1	How to use this interface manual	1
1.2	General interface description	1
1.3	Physical interface characteristics	3
1.4	Summary of SCSI commands and messages	4
2.0	SCSI bus	7
2.1	SCSI bus signals	8
2.1.1	Drive select	9
2.1.2	Signal values	9
2.1.3	OR-Tied signals	9
2.1.4	Signal sources	9
2.2	SCSI bus timing	10
2.2.1	Arbitration delay (2.4 μ s)	10
2.2.2	Assertion period (90 ns)	10
2.2.3	Bus Clear delay (800 ns)	11
2.2.4	Bus Free delay (800 ns)	11
2.2.5	Bus Set delay (1.8 μ s)	11
2.2.6	Bus Settle delay (400 ns)	11
2.2.7	Cable Skew delay (10 ns)	11
2.2.8	Data Release delay (400 ns)	11
2.2.9	Deskew delay (45 ns)	11
2.2.10	Disconnection delay (200 μ s)	11
2.2.11	Hold time (45 ns)	11
2.2.12	Negation period (90 ns)	11
2.2.13	Reset Hold time (25 μ s)	12
2.2.14	Selection Abort time (200 μ s)	12
2.2.15	Selection Timeout delay (250 ms)	12
2.2.16	Transfer Period (negotiated by SDTR message)	12
2.3	Fast synchronous transfer option timing	12
2.3.1	Fast Assertion period (30 ns)	12
2.3.2	Fast Cable Skew delay (5 ns)	12
2.3.3	Fast Deskew delay (20 ns)	12
2.3.4	Fast Hold time (10 ns)	12
2.3.5	Fast Negation period (30 ns)	12
3.0	Logical characteristics	13
3.1	SCSI bus phases	13
3.1.1	Bus Free phase	13
3.1.2	Arbitration phase	14
3.1.3	Selection phase	14
3.1.4	Reselection phase	16
3.1.5	Information Transfer phases	17
3.1.6	Command phase	19
3.1.7	Data phase	19
3.1.8	Status phase (see Section 4.3)	19
3.1.9	Message phase	20
3.1.10	Signal Restrictions Between phase	20

Table of contents *continued*

3.2	SCSI bus conditions	21
3.2.1	Attention conditions (see Unit Attention Condition, 4.6)	21
3.2.2	Reset condition	22
3.2.3	Contingent Allegiance condition	22
3.3	SCSI bus phase sequences	22
3.3.1	Nonarbitrating system	22
3.3.2	Arbitrating systems	22
3.3.3	All systems	23
3.4	SCSI pointers	24
3.4.1	Current pointers	25
3.4.2	Saved pointers	25
3.5	Message system specification	25
3.5.1	General message protocol	26
3.5.2	Messages - general	26
3.5.3	Message details	28
3.5.3.1	One and Two Byte message	28
3.5.3.2	Extended message (01h)	36
3.6	Message exception conditions handling	42
3.7	S.M.A.R.T. system	45
4.0	SCSI commands	47
4.1	Command implementation requirements	47
4.1.1	Reserved	47
4.2	Command descriptor block (CDB)	47
4.2.1	Operation code	48
4.2.2	Logical Unit Number (LUN)	49
4.2.3	Logical Block Address	50
4.2.4	Relative Address Bit	50
4.2.5	Transfer Length	50
4.2.6	Control Byte	51
4.3	Status	51
4.4	Command examples	53
4.4.1	Single Command example	53
4.4.2	Disconnect example	55
4.5	Timing examples	57
4.6	Unit attention condition	72
4.7	Queued I/O processes	73
4.7.1	Untagged queuing	73
4.7.2	Tagged queuing	73
4.8	Parameter Rounding	75
4.9	Command Processing considerations and exception conditions	75
4.9.1	Programmable Operating definition	75
4.9.2	Incorrect initiator connection	76
4.9.3	Selection of an Invalid Logical unit	76
5.0	Command descriptions	77
5.1	Command descriptions for all device types	77
5.1.1	Group 0 commands for all device types	77
5.1.1.1	Test Unit Ready (00h)	77
5.1.1.2	Request Sense command (03h)	78
5.1.1.2.1	Deferred errors	86
5.1.1.3	Inquiry command (12h)	87
5.1.1.3.1	Vital product data pages	91
5.1.1.4	Copy command (18h)	100
5.1.1.5	Receive Diagnostic results command (1Ch)	100
5.1.1.6	Send Diagnostic command (1Dh)	105

Table of contents *continued*

5.1.2	Group 1 commands for all device types	108
5.1.2.1	Compare command	(39h) 108
5.1.2.2	Copy and verify command	(3Ah) 108
5.1.2.3	Write Buffer command	(3Bh) 108
5.1.2.3.1	Combined Header and Data mode	(000b) 109
5.1.2.3.2	Write Data Only mode	(010b) 110
5.1.2.3.3	Download Microcode mode	(100b) 110
5.1.2.3.4	Download and Save Microcode mode	(101b) 110
5.1.2.3.5	Download Microcode with offsets	(110b) 111
5.1.2.3.6	Download Microcode with offsets and save mode	(111b) 112
5.1.2.4	Read Buffer command	(3Ch) 113
5.1.2.4.1	Read Combined Descriptor Header and Data mode	(000b) 113
5.1.2.4.2	Read data	114
5.1.2.4.3	Read Buffer Descriptor mode	(011b) 114
5.1.3	Group 2 Commands for All Device Types	115
5.1.3.1	Change Definition command.....	(40h) 115
5.1.3.2	Log Select command	(4Ch) 117
5.1.3.2.1	Write, Read and Verify Error Counter pages (code 02, 03, 05h)	123
5.1.3.2.2	Non-Medium Error page	(code 06h) 125
5.1.3.2.3	Cache Statistics page	(code 37h) 125
5.1.3.3	Log Sense command	(4Dh) 126
5.1.3.4	Other Group 2 commands	128
5.1.4	Groups 3 through 5 commands for all device types.....	128
5.1.5	Group 6 commands for all device types	128
5.1.6	Group 7 commands for all device types	128
5.2	Command descriptions for direct access devices.....	129
5.2.1	Group 0 commands for direct access devices	129
5.2.1.1	Rezero Unit command	(01h) 129
5.2.1.2	Format Unit command	(04h) 130
5.2.1.3	Reassign Blocks command	(07h). 137
5.2.1.4	Read command	(08h) 139
5.2.1.5	Write command	(0Ah) 141
5.2.1.6	Seek command	(0Bh) 142
5.2.1.7	Mode Select (6) command.....	(15h) 143
5.2.1.8	Reserve command.....	(16h) 147
5.2.1.9	Release command	(17h). 149
5.2.1.10	Mode Sense (6) command	(1Ah) 150
5.2.1.11	Start Stop Unit command	(1Bh) 183
5.2.1.12	Prevent/Allow Medium Removal command	(1Eh). 184
5.2.2	Group 1 command for direct access devices	184
5.2.2.1	Read Capacity command	(25h) 184
5.2.2.2	Read Extended command	(28h) 186
5.2.2.3	Write Extended command	(2Ah) 188
5.2.2.4	Seek Extended command.....	(2Bh) 189
5.2.2.5	Write and Verify command.....	(2Eh) 190
5.2.2.6	Verify command	(2Fh) 191
5.2.2.7	Synchronize Cache command	(35h) 192
5.2.2.8	Read Defect Data command	(37h) 193

Table of contents *continued*

5.2.2.9	Read Long command	(3Eh)	195
5.2.2.10	Write Long command	(3Fh)	197
5.2.2.11	Write Same command	(41h)	198
5.2.2.12	Reserve (10) command	(56h)	199
5.2.2.13	Release (10) command	(57h)	200
5.2.2.14	Mode Select (10) command	(55h)	201
5.2.2.15	Mode Sense (10) command	(5Ah)	202
5.2.2.16	Pre-Fetch command	(34h)	203
5.3	Exclusive Or commands		204
5.3.1	XDWRITE command	(50h)	205
5.3.2	XPWRITE command	(51h)	206
5.3.3	XDREAD command	(52h)	207
5.3.4	XDWRITE EXTENDED command	(80h)	208
5.3.5	REBUILD command	(81h)	210
5.3.6	REGENERATE command	(82h)	213
6.0	Seagate Technical Support Services		215
A.1.0	Appendix A. SCSI Configured Auto Magically (SCAM).....		A-1

List of Figures

Figure 2.0-1	SCSI ID bits	7
Figure 3.3-1	Phase sequences with arbitration	23
Figure 3.3-2	Phase sequences without arbitration	23
Figure 3.4-1	Simplified SCSI System	24
Figure 3.5.3-1	Wide SCSI byte ordering	41
Figure 4.4-1.	Single command example	54
Figure 4.4-2.	Disconnect example	56

1.0 Interface requirements

1.1 How to use this interface manual

This specification is designed to provide a universal detailed description of the SCSI interface for those drive products whose Product Manuals (Volume 1) do not contain the details of how the SCSI interface is implemented by that drive.

Note: Volume 1 Product Manuals have tables that specify which SCSI-1 or SCSI-2/SCSI-3 features they implement, what the default parameters are for the various features they implement and which parameters are changeable and which are not.

No method exists at present to inform an initiator if a target supports “SCSI-3” features as opposed to only SCSI-2 features. A few “SCSI-3” features are supported by Seagate drives, but no attempt has been made herein to differentiate between SCSI-2 and “SCSI-3” features. Therefore, when an Inquiry command reports what the ANSI approved version of the drive is, it reports either SCSI-1 or SCSI-2, where “SCSI-2” means SCSI-2 features plus some “SCSI-3” features.

No attempt is made in this universal specification to specify which descriptions or tables apply to SCSI-1 and which to SCSI-2 or SCSI-3. The combination of this general specification with the details in the individual drive Product Manuals (Volume 1) provides a description of the individual drive implementation of the SCSI interface.

This interface manual is not intended to be stand-alone text on SCSI-1 or SCSI-2/SCSI-3 features. Reference must be made back to the individual drive Product Manuals to find out what are SCSI-1 and what are SCSI-2/SCSI-3 features.

This specification is Volume 2 of a set of manuals that is made up of separate drive Product Manuals (Volume 1) and this manual. This Volume 2 Manual is referenced by other Volume 1 Product Manuals representing the drives listed below.

Product Manuals for the following models reference this volume: ST12400, ST11900, ST31200, ST11950, ST12450, ST12550, ST15150, ST3655, ST3390, ST3285, ST31230, ST31231, ST32430, ST15230, ST31250, ST32151N, ST31051, ST32550, ST32155, ST31055, ST34371, ST32171, ST410800, ST410800, ST18771, ST19171, ST34501, ST19101, ST423451, ST52160, ST118273, ST42345, ST39173, ST34573, ST34572, ST32272, ST39102, and ST118202.

1.2 General interface description

This Product Manual describes the Seagate Technology, Inc. subset of the SCSI (Small Computer Systems Interface) as implemented on the Seagate built drives listed above. The interface is compatible with the SCSI Interface Specifications of the ANSI SCSI-1 standard, the ANSI SCSI-2 Standard and the common command set (CCS) document, Revision 4.B. The drives covered by this Product Manual are classified as “Intelligent” peripherals.

The Seagate SCSI interface described herein consists of a 9 or 18 bit bidirectional bus (8 data + 1 parity or 16 data + 2 parity) plus 9 control signals supporting multiple initiators, disconnect/ reconnect, self configuring host software, automatic features that relieve the host from the necessity of knowing the physical architecture of the target (logical block addressing is used), and some other miscellaneous features.

The SCSI physical interface uses either single ended drivers and receivers or differential drivers and receivers and uses asynchronous or synchronous communication protocols. The bus interface transfer rate for asynchronous or synchronous is given in individual drive Volume 1 Product Manuals. The bus protocol supports multiple initiators, disconnect/reconnect, additional messages plus 6 byte and 10 byte Command Descriptor Blocks.

Unless specified otherwise in the individual drive Product Manuals (Vol. 1), the drive is always a target, and never an initiator. For certain commands, which may or may not be supported by a particular drive model, the drive must act as an initiator, but does not otherwise do so. For purposes of this specification, "drive" may be substituted for the word "target" wherever "target" appears.

GLOSSARY

Arbitration - SCSI bus phase wherein SCSI devices try to gain control of the SCSI bus to operate as an initiator or target (see Section 3.1.2).

Byte - This term indicates an 8 bit hexadecimal construction.

Command Descriptor Block (CDB) - The structure used to communicate requests from an initiator to a target.

Connect - The function that occurs when an initiator selects a target to start an operation.

Disconnect - The function that occurs when a target releases control of the SCSI bus, allowing it to go to the Bus Free phase.

FRU (Field Replaceable Unit) - An assembly that is believed faulty based on test results. A value of 00h indicates an unknown cause or the end of a list of known possible causes. Nonzero values have product unique meanings.

Initiator - A SCSI device (usually a host system) that requests an operation to be performed by another SCSI device.

Intermediate Status - A status code sent from a target to an initiator upon completion of each command, except the last command, in a set of linked commands.

I/O Process - An I/O process consists of one initial connection and zero or more reconnections, all pertaining to a single command or group of linked commands. More specifically, the connection(s) pertain to a nexus as defined below in which one or more command descriptor blocks are usually transferred. An I/O process begins with the establishment of a nexus. An I/O process normally ends with the BUS Free phase following successful transfer of a COMMAND COMPLETE, ABORT, ABORT TAG, or CLEAR QUEUE message. An I/O process also ends when a hard RESET condition occurs, an unexpected BUS FREE phase occurs, or when the BUS FREE phase occurs following a BUS DEVICE RESET message.

I T nexus - A nexus prior to the successful receipt of an IDENTIFY message, at which time the nexus is changed to an I T L nexus. (See glossary word "Nexus").

I T L nexus - A nexus that exists between an initiator and a Logical Unit. This relationship replaces the prior I T nexus. (See glossary word "Nexus").

I T L Q nexus - A nexus between an initiator, a Logical Unit, and a queue tag following the successful receipt of one of the QUEUE messages. This relationship replaces the prior I T L nexus. (See glossary word "Nexus").

Logical Unit - A physical device or virtual device addressable through a target. The drive is a target but also a Logical Unit.

Logical Unit Number - An encoded three bit identifier for the logical unit. The drive is considered Logical Unit number zero.

LSB - Least significant byte

MSB - Most significant byte

ms - millisecond

LUN - Logical unit number

mm - Millimeter

nexus - A relationship that begins with the establishment of an initial connection and ends with the completion of the I/O process. The relationship starts as an association between the initiator and a selected target. The relationship may be restricted to specify a single logical unit or target routine by the successful transfer of an IDENTIFY message. The relationship may be further restricted by the successful transfer of a queue tag message.

ns - Nanosecond

One - A logical true signal value, (assertion).

Page - Several commands use regular parameter structures that are referred to as pages. These pages are identified with a value known as a page code.

Queue - This term refers to the command queue used in tagged queuing (see 4.7.2).

Queue Tag - The value associated with an I/O process that uniquely identifies it from other queued I/O processes on the logical unit for the same initiator.

Reconnect - The function that occurs when a target selects an initiator to continue an operation after a disconnect.

Reconnection - A reconnection exists from the assertion of the BSY signal in a RESELECTION phase until the next BUS FREE phase occurs. A reconnection can only occur between a target and an initiator.

Reserved - The term used for bits, bytes, fields, and code values that are set aside for future standardization.

SCSI Address - The octal representation of the unique address (0-7) assigned to a SCSI device. This address would normally be assigned and set in the SCSI device during system installation (see individual drive Product Manuals).

SCSI ID - The bit significant representation of the SCSI address referring to one of the signal lines DB(7-0).

SCSI device - A host computer adapter or a peripheral controller or an intelligent peripheral that can be attached to the SCSI bus.

Signal Assertion - The act of driving a signal to the true state.

Signal Negation - The act of driving a signal to the false state or allowing the cable terminators to bias the signal to the false state (by placing the driver in the high impedance condition).

Signal Release - The act of allowing the cable terminators to bias the signal to the false state (by placing the driver in the high impedance condition).

xxh - Numbers followed by lower case h are hexadecimal values. All other numbers are decimal values.

Status - One byte of information sent from a target to an initiator upon completion of each command.

Target - A SCSI device that performs an operation requested by an initiator.

μs - Microsecond.

Vendor Unique - In this specification, this term indicates bits, fields, or code values that are vendor specific.

Zero - A logical false signal value, (negation).

1.3 Physical interface characteristics

The physical interface characteristics (cables, connectors, electrical descriptions, termination requirements, etc.) for the drives covered by this Interface Manual are found in each individual Product Manual, since these features are not the same for all drives.

1.4 Summary of SCSI commands and messages

Following is an alphabetical table listing the SCSI commands described in this manual. Details are given in Section 5.

Command name	Hex code	Device type	Page number
Change Definition	40	All	115
Format Unit	04	dir. access	130
Inquiry	12	All	87
Log Select	4C	All	117
Log Sense	4D	All	126
Mode Select (6 byte)	15	dir. access	143
Mode Select (10 byte)	55	dir. access	201
Mode Sense (6 byte)	1A	dir. access	150
Mode Sense (10 byte)	5A	dir. access	202
Prefetch	34	dir. access	203
Read	08	dir. access	139
Read Buffer	3C	All	113
Read Capacity	25	dir. access	184
Read Defect Data	37	dir. access	193
Read Extended	28	dir. access	186
Read Long	3E	dir. access	195
Reassign Blocks	07	dir. access	137
Rebuild	81	dir. access	210
Receive Diagnostic Results	1C	All	100
Regenerate	82	dir. access	213
Release (6 byte)	17	dir. access	149
Release (10 byte)	57	dir. access	200
Request Sense	03	All	78
Reserve (6 byte)	16	dir. access	147
Reserve (10 byte)	56	dir. access	199
Rezero	01	dir. access	129
Seek	0B	dir. access	142
Seek Extended	2B	dir. access	189
Send Diagnostics	1D	all	105
Start/Stop Unit	1B	dir. access	183
Synchronize Cache	35	dir. access	192
Test Unit Ready	00	All	77
Verify	2F	dir. access	191
Write	0A	dir. access	141
Write and Verify	2E	dir. access	190
Write Buffer	3B	All	108
Write Extended	2A	dir. access	188
Write Long	3F	dir. access	197
Write Same	41	All	198
XD Read	52	dir. access	207
XD Write	50	dir. access	205
XD Write Extended	80	dir. access	208
XP Write	51	dir. access	206

Following is an alphabetical summary of the SCSI messages described in this manual. Details are given in Section 3.5.

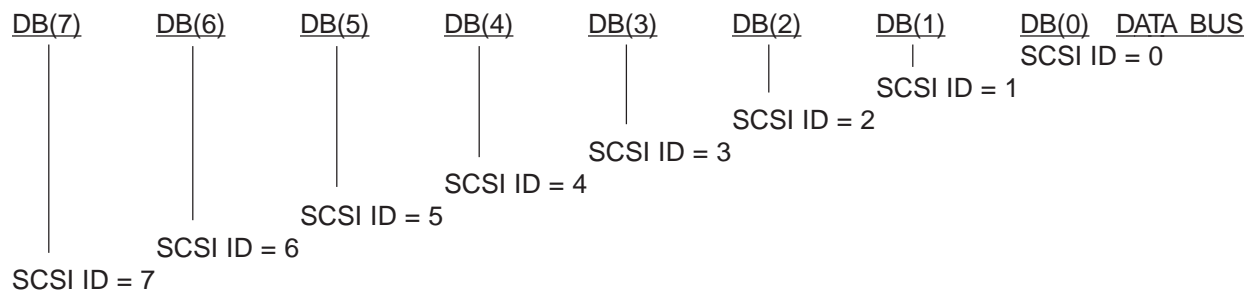
Message Name	Hex Code	Page number
Abort	06	28
Abort Tag	0D	29
Bus Device Reset	0C	29
Clear Queue	0E	29
Command Complete	00	29
Continue I/O Process	12	29
Disconnect	04	30
Extended Message	01	36
Identify	80-FF	30
Ignore Wide Residue	23	42
Initiate Recovery	0F	not supported
Initiator Detected Error	05	31
Linked Command Complete	0A	31
Linked Command Complete (with flag)	0B	31
Message Parity Error	09	31
Message Reject	07	32
Modify Data Pointer	01	32, 36 (extended message)
No Operation	08	32
Queue Tag Messages		33
Head of Queue Tag	21	33
Ordered Queue Tag	22	34
Simple Queue Tag	20	34
Release Recovery	10	not supported
Restore Pointers	03	34
Save Data Pointers	02	34
Synchronous Data Transfer Request	01	34, 36 (extended message)
Terminate I/O Process	11	34 (optional)
Target Transfer Disable	13	35
Wide Data Transfer Request	01	39 (extended message)

2.0 SCSI bus

This manual discusses only the “logical” and timing characteristics of the SCSI system and interface. The SCSI bus physical characteristics (voltages, connector configurations, pinouts, etc.) are given in the individual drive Product Manuals (Volume 1) Section “Interface requirements”, which covers all of the interface requirements and SCSI features supported by the drive described in the particular Product Manual being referenced.

Communication on the SCSI Bus is allowed between only two SCSI devices at a time. Some Seagate drives support systems with a maximum of eight SCSI devices including the host computer(s) connected to the SCSI bus. Some Seagate drives support systems with a maximum of sixteen SCSI devices on the SCSI bus. Each SCSI device has a SCSI ID Bit assigned as shown in Figure 2.0-1. The SCSI ID is assigned by installing from 0 to 3 (8 device systems) jumper plugs or 0-4 (16 device systems) jumper plugs onto a connector in a binary coded configuration during system configuration. Some drive models have an interface that includes the SCSI bus ID lines, so that the host can set the drive ID over the interface. See individual drive Product Manual, Section “Option/configuration headers”.

When two SCSI devices communicate on the SCSI Bus one acts as an initiator and the other acts as a target. The initiator (typically a host computer) originates an operation and the target performs the operation. The drive always operates as a target, unless specified otherwise (i.e., certain commands are supported) in the individual drive Product Manual.



Additional SCSI ID bits for devices that support 16 devices on the SCSI bus.

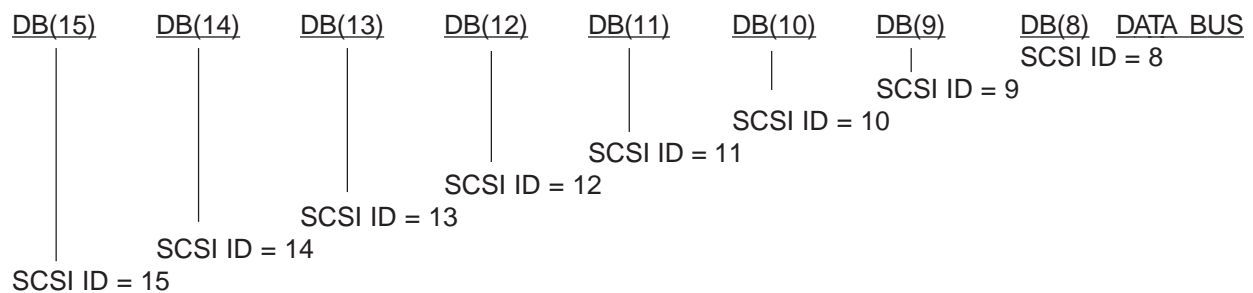


Figure 2.0-1. SCSI ID bits

The Host Adapter/Initiator must be identified by one of the eight SCSI Device Addresses. Make sure that none of the devices on the SCSI bus have duplicate addresses.

Certain SCSI bus functions are assigned to the initiator and certain SCSI bus functions are assigned to the target. The initiator will select a particular target. The target will request the transfer of Command, Data, Status or other information on the data bus.

Information transfers on the data bus are interlocked and follow a defined REQ/ACK Handshake protocol. One byte of information will be transferred with each handshake. Synchronous data transfers do not require a one for one interlocking of REQ/ACK signals, but the total number of REQ pulses in a particular data transfer event must equal the total number of ACK pulses. Synchronous data transfer option is described in Paragraph 3.5.3.2 and 3.1.5.2.

The drive supports single initiator, single target; single initiator, multiple target; multiple initiator, single target; or multiple initiator, multiple target bus configurations.

2.1 SCSI bus signals

There are ten control and eighteen data signals, as listed below:

- BSY
- C/D
- MSG
- DIFFSENS
- SEL
- I/O
- REQ
- DB(7-0, P); DB(15-8,P1)
- ACK
- ATN
- RST

Some drive models have a single 80 pin I/O connector that contains additional interface lines that carry drive configuration select signals. These are peculiar to certain drives and are not SCSI standard signals. These are described in the drive model's Volume 1 Product manual, but not here.

The 28 SCSI standard signals are described as follows:

BSY (Busy) - An "OR-tied" signal to indicate the bus is being used.

SEL (Select) - A signal used by an initiator to select a target, or by a target to reselect an initiator.

C/D (Control/Data) - A signal driven by a target to indicate whether Control or Data information is on the Data Bus. Assertion (see Paragraph 2.1.2) indicates Control.

I/O (Input/Output) - A signal driven by a target to control the direction of data movement on the Data Bus with respect to an initiator. Assertion indicates input to the initiator. This signal also distinguishes between Selection and Reselection phases.

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 REQ/ACK data transfer handshake.

ACK (Acknowledge) - A signal driven by an initiator to indicate an acknowledgment for a REQ/ACK data transfer handshake.

ATN (Attention) - A signal driven by an initiator to indicate the Attention condition. It is used to request to send a message out to the target. See paragraph 3.2.1. If an initiator asserts ATN while asserting SEL it indicates to the target that the initiator supports messages other than command complete.

RST (Reset) - An "OR-tied" signal that indicates the Reset condition.

DIFFSENS (Differential Sense) - When the drive has differential SCSI I/O circuits, the DIFFSENS signal disables the drive's differential driver/receiver circuits if the SCSI I/O cable is plugged in upside down, or if a single-ended SCSI I/O cable is plugged into a differential I/O drive. Disabling the differential I/O drivers/receivers is necessary to prevent burning them out if a grounded I/O line is connected to any of the differential circuit outputs, which are at a positive voltage (+2 V or +3 V) when not disabled.

DB(7-0,P) and DB(15-8,P1) (Data Bus) - Sixteen data bit signals, plus parity bit signals form a Data Bus. DB(7) is the most significant bit and has the highest priority during the Arbitration phase (on both eight and sixteen device systems). Bit number significance, and priority decrease downward to DB(0), and then from DB15 down to DB8 (DB0 is higher than DB15). A data bit is defined as one when the signal is asserted and is defined as zero when the signal is negated.

Data parity DB(P) and DB(P1) is odd - The use of parity is a system option. The drive always checks parity on the data bits, but has the capability to enable/disable parity error reporting to the host. See configuration selection in the applicable Product Manual. Parity checking is not valid during the Arbitration phase.

Greater detail on each of the SCSI Bus signals is found in the following sections.

2.1.1 Drive Select

For SCSI ID selection install drive select jumpers as shown in configuration selection figure in applicable Product Manual. Refer to the individual drive Product Manual for the location of the drive select header. The drive using the eight bit data interface can have one of eight ID bits selected by installing 0 to 3 jumpers in a binary coded configuration on the drive select header. Drives using the 16 bit data interface can have one of sixteen ID bits selected by installing 0 to 4 jumpers in a binary coded configuration on the drive select header.

2.1.2 Signal Values

Signals may assume true or false values. There are two methods of driving these signals. In both cases, the signal shall be actively driven true, or asserted. In the case of OR-tied drivers, the driver does not drive the signal to the false state, rather the bias circuitry of the bus terminators pulls the signal false whenever it is released by the drivers at every SCSI device. If any driver is asserted, then the signal is true. In the case of non-OR-tied drivers, the signal may be negated. Negated means that the signal may be actively driven false, or may be simply released (in which case the bias circuitry pulls it false), at the option of the implementor.

2.1.3 OR-Tied signals

The BSY and RST signals shall be OR-tied only. In the ordinary operation of the bus, these signals are simultaneously driven true by several drivers. No signals other than BSY, RST, and DB(P) are simultaneously driven by two or more drivers, and any signal other than BSY and RST may employ OR-tied or non-OR-tied drivers. DB(P) shall not be driven false during the Arbitration phase. There is no operational problem in mixing OR-tied and non-OR-tied drivers on signals other than BSY and RST.

2.1.4 Signal sources

Table 2.1.4-1 indicates which type of SCSI device is allowed to source each signal. All SCSI device drivers that are not active sources shall be in the passive state. Note that the RST signal may be sourced by any SCSI device at any time. The drive functions as a target.

2.2 SCSI bus timing

Unless otherwise indicated, the delay time measurements for each SCSI device, defined in Paragraphs 2.2.1 through 2.2.14, shall be calculated from signal conditions existing at that SCSI device's own SCSI bus connection. Thus, these measurements (except skew delay) can be made without considering delays in the cable.

See Section 2.3 for Fast Synchronous transfer option timing.

Table 2.1.4-1. Signal sources

Bus Phase	BSY	Signals SEL	C/D, I/O, MSG, REQ	ACK/ATN	DB(7-0,P) (15-8,P1)
Bus Free	None	None	None	None	None
Arbitration	All	Winner	None	None	SCSI ID
Selection	I&T	Init.	None	Init.	Init.
Reselection	I&T	Target	Target	Init.	Target
Command	} Target	None	Target	Init.	Init.
Data In		None	Target	Init.	Target
Data Out		None	Target	Init.	Init.
Status		None	Target	Init.	Target
Message In		None	Target	Init.	Target
Message Out		None	Target	Init.	Init.

ALL: The signal shall be driven by all actively arbitrating SCSI devices.

SCSI ID: A unique data bit (the SCSI ID) shall be driven by each actively arbitrating SCSI device: the other seven data bits shall be released (i.e., not driven) by this SCSI device. The parity bit [DB(P), DB(P1)] may be undriven or driven to the true state, but shall never be driven to the false state during this phase.

I&T: The signal shall be driven by the initiator, target, or both, as specified in the Selection phase and Reselection phase.

Init: If this signal is driven, it shall be driven only by the active initiator.

None: The signal shall be released; that is, not be driven by any SCSI device. The bias circuitry of the bus terminators pulls the signal to the false state.

Winner: The signal shall be driven by the one SCSI device that wins arbitration.

Target: If the signal is driven, it shall be driven only by the active target.

2.2.1 Arbitration delay (2.4 μ s)

The minimum time a SCSI device shall wait from asserting BSY for arbitration until the Data Bus can be examined to see if arbitration has been won. There is no maximum time.

2.2.2 Assertion period (90 ns)

The minimum time that a target shall assert REQ while using synchronous data transfers. Also, the minimum time that an initiator shall assert ACK while using synchronous data transfers.

2.2.3 Bus clear delay (800 ns)

The maximum time for a SCSI device to stop driving all bus signals after:

- (1) The Bus Free phase is detected (BSY and SEL both negated for a bus settle delay).
- (2) SEL is received from another SCSI device during the Arbitration phase.
- (3) The transition of RST to assertion.

Note. For the first condition above, the maximum time for a SCSI device to clear the bus is 1200 ns from BSY and SEL first becoming both negated. If a SCSI device requires more than a bus settle delay to detect Bus Free phase, it shall clear the bus within a Bus Clear delay minus the excess time.

2.2.4 Bus free delay (800 ns)

The minimum time that a SCSI device shall wait from its detection of the Bus Free phase (BSY and SEL both negated for a bus settle delay) until its assertion of BSY when going to the Arbitration phase.

2.2.5 Bus set delay (1.8 μ s)

The maximum time for a SCSI device to assert BSY and its SCSI ID bit on the Data Bus after it detects Bus Free phase (BSY and SEL both negated for a bus settle delay) for the purpose of entering the Arbitration phase.

2.2.6 Bus settle delay (400 ns)

The time to wait for the bus to settle after changing certain control signals as specified in the protocol definitions.

2.2.7 Cable skew delay (10 ns)

The maximum difference in propagation time allowed between any two SCSI bus signals when measured between any two SCSI devices.

2.2.8 Data release delay (400 ns)

The maximum time for an initiator to release the Data Bus signals following the transition of the I/O signal from negation to assertion.

2.2.9 Deskew delay (45 ns)

The minimum time required for deskew of certain signals.

2.2.10 Disconnection delay (200 μ s)

The minimum time that a target shall wait after releasing BSY before participating in an Arbitration phase when honoring a Disconnect message from the initiator.

2.2.11 Hold time (45 ns)

The 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 target, respectively, while using synchronous data transfers.

2.2.12 Negation period (90 ns)

The minimum time that a target shall negate REQ while using synchronous data transfers. Also, the minimum time that an initiator shall negate ACK while using synchronous data transfers.

2.2.13 Reset hold time (25 μ s)

The minimum time for which RST is asserted. There is no maximum time.

2.2.14 Selection abort time (200 μ s)

The maximum time that a target (or initiator) shall take from its most recent detection of being selected (or reselected) until asserting a BSY response. This timeout is required to ensure that a target (or initiator) does not assert BSY after a Selection (or Reselection) phase has been aborted. This is not the selection timeout period; see Sections 3.1.3.5 and 3.1.4.2 for a complete description.

2.2.15 Selection timeout delay (250 ms recommended)

The minimum time an initiator (or target) should wait for a BSY response during the Selection (or Reselection) phase before starting the timeout procedure. The drive implements this 250 ms selection timeout delay.

2.2.16 Transfer period (negotiated by Synchronous Data Transfer message)

The minimum time allowed between the leading edges of successive REQ pulses and of successive ACK pulses while using synchronous data transfers. (See Sections 3.1.5.2 and 3.5.3.2)

2.3 Fast synchronous transfer option

When devices negotiate a synchronous data transfer period of less than 200 ns they are said to be using "fast synchronous data transfers". Devices which negotiate a synchronous data transfer period greater than or equal to 200 ns use timing parameters specified in 2.2. When a fast synchronous data transfer period is negotiated, those specific times redefined in this section are used; those not redefined remain the same. The minimum synchronous data transfer period is 100 ns.

2.3.1 Fast Assertion period (30 ns)

This value is the minimum time that a target shall assert REQ while using fast synchronous data transfers. Also, the minimum time that an initiator shall assert ACK while using fast synchronous data transfers.

2.3.2 Fast cable skew delay (5 ns)

This value is the maximum difference in propagation time allowed between any two SCSI bus signals measured between any two SCSI devices while using fast synchronous data transfers.

2.3.3 Fast deskew delay (20 ns)

This value is the minimum time required for deskew of certain signals while using fast synchronous data transfers.

2.3.4 Fast hold time (10 ns)

This value is the 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 target respectively, while using fast synchronous data transfers.

2.3.5 Fast negation period (30 ns)

This value is the minimum time that a target shall negate REQ while using fast synchronous data transfers. Also, the minimum time that an initiator shall negate ACK while using fast synchronous data transfers.

3.0 Logical characteristics

The operations of the SCSI bus as described in Section 3 are supported by the drive, as specified in each drive's Product Manual (Vol. 1). The drive always functions as the target unless otherwise stated.

3.1 SCSI bus phases

The drive responds to 8 distinct bus phases.

Bus Free phase
Arbitration phase
Selection phase
Reselection phase

Command Phase
Data (in and out)

Status (in only)
Message (in and out)

} These phases are collectively termed the Information transfer phases

The SCSI Bus can never be in more than one phase at a time.

3.1.1 Bus free phase

The Bus Free phase indicates that no SCSI device is actively using the SCSI bus and it is available for subsequent users. In some cases a target reverts to the Bus Free phase to indicate an error condition that it has no other way to handle. This is called an unexpected disconnect.

SCSI devices shall detect the Bus Free phase after SEL and BSY are both false for at least a bus settle delay.

SCSI devices shall release all SCSI bus signals within a bus clear delay after BSY and SEL are continuously negated for a bus settle delay. If a SCSI device requires more than a bus settle delay to detect the Bus Free phase, it shall release all SCSI bus signals within a bus clear delay minus the excess time to detect the Bus Free phase. The total time to clear the SCSI bus shall not exceed a bus settle delay plus a bus clear delay.

Initiators normally do not expect BUS FREE phase to begin because of the target's release of the BSY signal except after one of the following occurrences:

- (1) after a reset condition is detected.
- (2) after an ABORT message is successfully received by a target.
- (3) after a BUS DEVICE RESET message is successfully received by a target.
- (4) after a DISCONNECT message is successfully transmitted from a target (see 3.5.3.1).
- (5) after a COMMAND COMPLETE message is successfully transmitted from a target (see 3.5.3.1).
- (6) after a RELEASE RECOVERY message is successfully received by a target.
- (7) after an ABORT TAG message is successfully received by a target.
- (8) after a CLEAR QUEUE message is successfully received by a target.

The BUS FREE phase may also be entered after an unsuccessful selection or reselection, although in this case it is the release of the SEL signal rather than the release of the BSY signal that first establishes the BUS FREE phase.

If an initiator detects the release of the BSY signal by the target at any other time, the target is indicating an error condition to the initiator. The target may perform this transition to the BUS FREE phase independent of the state of the ATN signal. The initiator shall manage this condition as an unsuccessful I/O process termination. The target terminates the I/O process by clearing all pending data and status information for the affected logical unit or target routine. The target may optionally prepare sense data that may be retrieved by a REQUEST SENSE command. When an initiator detects an unexpected disconnect, it is recommended that a REQUEST SENSE command be attempted to obtain any valid sense data that may be available.

3.1.2 Arbitration phase

The Arbitration phase allows one SCSI device to gain control of the SCSI bus so that it can assume the role of an initiator or target. The drive arbitrates for the bus as a target implementing reselection or when performing AEN (if AEN is implemented). The drive supports arbitration by multiple SCSI devices.

The procedure for a SCSI device to obtain control of the SCSI bus is as follows:

1. The SCSI device shall first wait for the Bus Free phase to occur. The Bus Free phase is detected when BSY and SEL are simultaneously and continuously negated for a minimum of a bus settle delay. (Implementors note: This bus settle delay is necessary because a transmission line phenomenon known as a "Wire-OR glitch" may cause BSY to briefly appear negated, even though it is being asserted.)
2. The SCSI device shall wait a minimum of a bus free delay after detection of the Bus Free phase (i.e. after BSY and SEL are both negated for a bus settle delay) before driving any signal.
3. Following the bus free delay in Step (2), the SCSI device may arbitrate for the SCSI bus by asserting both BSY and its own SCSI ID, however the SCSI device shall not arbitrate (i.e. assert BSY and its SCSI ID) if more than a bus settle delay has passed since the Bus Free phase was last observed. (Implementors Note: There is no maximum delay before asserting BSY and the SCSI ID following the bus free delay in Step (2) as long as the bus remains in the Bus Free phase. However, SCSI devices that delay longer than a bus settle delay plus a bus set delay from the time when BSY and SEL are first negated may fail to participate in arbitration when competing with faster SCSI devices.)
4. After waiting at least an arbitration delay (measured from its assertion of BSY) the SCSI device shall examine the Data Bus. If a higher priority SCSI ID bit is true on the Data Bus [DB(7) is the highest], the SCSI device has lost the arbitration and the SCSI device must release its signals and return to Step (1). If no higher priority SCSI ID bit is true on the Data Bus, the SCSI device has won the arbitration and it shall assert SEL. Any other SCSI device that is participating in the Arbitration phase has lost the arbitration and shall release BSY and its SCSI ID bit within a bus clear delay after SEL becomes true. A SCSI device that loses arbitration may return to Step (1).
5. The SCSI device that wins arbitration shall wait at least a bus clear delay plus a bus settle delay after asserting SEL before changing any signals.

Note. The SCSI ID bit is a single bit on the Data Bus that corresponds to the SCSI device's unique SCSI address. All other seven Data Bus bits shall be released by the SCSI device. Parity is not valid during the Arbitration phase, DB(P) may be undriven or driven to the true state, but shall not be driven to the false state.

3.1.3 Selection phase

The Selection phase allows an initiator to select a target for the purpose of initiating some target function (e.g., Read or Write command).

Note. During the Selection phase the I/O signal shall be negated so this phase can be distinguished from the Reselection phase.

3.1.3.1 Nonarbitrating system

In systems with the Arbitration phase not implemented, the initiator shall first detect the Bus Free phase and then wait a minimum of a bus clear delay. Then, except in certain single initiator environments with initiators employing the single initiator option (see 3.1.3.4), the initiator shall assert the desired target's SCSI ID and its own initiator SCSI ID on the Data Bus. After two deskew delays, the initiator shall assert SEL.

3.1.3.2 Arbitrating systems

In systems with the Arbitration phase implemented, the SCSI device that won the arbitration has both BSY and SEL asserted and has delayed at least a bus clear delay plus a bus settle delay before ending the Arbitration phase. The SCSI device that won the arbitration becomes an initiator by releasing I/O. Except in certain single initiator environments with initiators employing the single initiator option (see 3.1.3.4), the initiator shall set the Data Bus to a value which is the OR of its SCSI ID bit and the target's SCSI ID bit. The initiator shall then wait at least two deskew delays and release BSY. The initiator shall then wait at least a bus settle delay before looking for a response from the target.

3.1.3.3 All systems

In all systems, the target shall determine that it is selected when SEL and its SCSI ID bit are true and the BSY and I/O signals are false for at least a bus settle delay. The selected target will examine the Data Bus in order to determine the SCSI ID of the selecting initiator unless the initiator employed the single initiator option (see 3.1.3.4). The selected target shall then assert BSY within a selection abort time of its selection; this is required for correct operation of the timeout procedure. In systems with parity implemented, the target shall not respond to a selection if bad parity is detected. Also, if more than two SCSI ID bits are on the Data Bus, the target shall not respond to selection. At least two deskew delays after the initiator detects BSY is asserted, it shall release SEL and may change the Data Bus.

3.1.3.4 Single initiator option

Initiators that do not implement the Reselection phase, and do not operate in the multiple initiator environment, are allowed to set only the target's SCSI ID bit during the Selection phase. This makes it impossible for the target to determine the initiator's SCSI ID.

3.1.3.5 Selection time out procedure

A Selection timeout procedure is specified for clearing the SCSI bus. If the initiator waits a minimum of a selection timeout delay and there has been no BSY response from the target, the initiator shall continue asserting SEL and shall release the Data Bus. If the initiator has not detected BSY to be asserted after at least a selection abort time plus two deskew delays, the initiator shall release SEL allowing the SCSI bus to go to the Bus Free phase. SCSI devices shall ensure when responding to selection that the selection was still valid within a selection abort time of their assertion of BSY. Failure to comply with this requirement could result in an improper selection (two targets connected to the same initiator, wrong target connected to an initiator, or a target connected to no initiator).

The drive supports systems that implement this procedure.

3.1.4 Reselection phase

Reselection is a phase that allows a target to reconnect to an initiator for the purpose of continuing some operation that was previously started by the initiator but was suspended by the target (i.e., the target disconnected by allowing a Bus Free phase to occur before the operation was complete).

Reselection can be used only in systems that have Arbitration phase implemented.

The drive implements the Reselection phase if the system is capable of supporting Reselection.

ATN during Selection implies that the host supports messages other than command complete. Bit 6 of the Identify message indicates that the Disconnected/Reconnect privilege is granted.

ATN not asserted during Selection means that the host doesn't support the Identify message, so the Disconnected/Reconnect privilege cannot be granted. If a target is not granted disconnect/reconnect privileges, it shall not disconnect until the current command is completed. None of the Disconnect/Reconnect Control Page (page 02h) parameters of Mode Select are applicable if disconnect/reconnect privileges are not granted by the initiator. Operations taking longer than the bus inactivity limit specified do not cause a disconnect.

3.1.4.1 Reselection Procedure

Upon completing the Arbitration phase, the winning SCSI device has both BSY and SEL asserted and has delayed at least a bus clear delay plus a bus settle delay. The winning SCSI device becomes a target by asserting the I/O signal. That device shall also set the Data Bus to a value that is the OR of its SCSI ID bit and the initiator's SCSI ID bit. The target shall wait at least two deskew delays and release BSY. The target shall then wait at least a bus settle delay before looking for a response from the initiator.

The initiator shall determine that it is reselected when SEL, I/O, and its SCSI ID bit are true and BSY is false for at least a bus settle delay. The reselected initiator may examine the Data Bus to determine the SCSI ID of the reselecting target.

The reselected initiator shall then assert BSY within a selection abort time of its most recent detection of being reselected; this is required for correct operation of the timeout procedure. In systems with parity implemented, the initiator shall not respond to Reselection if bad parity is detected. The initiator shall not respond to a Reselection if more than two SCSI ID bits are on the Data Bus.

After the target detects BSY, it shall also assert BSY and wait at least two deskew delays and then release SEL. The target may then change the I/O signal and the Data Bus. After the reselected initiator detects SEL false, it shall release BSY. The target shall continue asserting BSY until the target is ready to relinquish the SCSI bus.

Note. When the target is asserting BSY, a transmission line phenomenon known as a "Wired-OR glitch" may cause BSY to appear false for up to a round trip propagation delay following the release of BSY by the initiator. This is the reason why the Bus Free phase is recognized only after both BSY and SEL are continuously false for a minimum of a bus settle delay. Cables longer than 25 meters should not be used even if the chosen driver, receiver, and cable provide adequate noise margins, because they increase the duration of the glitch and could cause SCSI devices to inadvertently detect the Bus Free phase.

3.1.4.2 Reselection timeout procedure

This Reselection timeout procedure is specified for clearing the SCSI bus during a Reselection phase. If the target waits a minimum of a selection timeout period and there has been no BSY response from the initiator, the target shall continue asserting SEL and I/O and shall release all Data Bus signals. If the target has not detected BSY to be true after at least a selection abort time plus two deskew delays, the target shall release SEL and I/O allowing the SCSI bus to go to the Bus Free phase. SCSI devices that respond to Reselection shall ensure that the Reselection was still valid within a selection abort time of their assertion of BSY. Failure to comply with this requirement could result in an improper Reselection (two initiators connected to the same target or the wrong initiator connected to a target).

If an initiator times out while waiting to be reselected, the initiator should attempt to select and issue Request Sense to determine if the previous command is:

1. Still in process (Busy Status is returned),
2. Aborted with valid Request Sense data, or
3. Aborted without valid Request Sense data.

3.1.5 Information transfer phases

NOTE: The Command, Data, Status, and Message phases are grouped together as information transfer phases because they are all used to transfer data or control information via the Data Bus. The actual contents of the information is beyond the scope of this section.

The C/D, I/O, and MSG signals are used to distinguish between the different information transfer phases. (See Table 3.1.5-1). The target drives these three signals and therefore controls all changes from one phase to another. The initiator can request a Message Out phase by asserting ATN, while the target can cause the Bus Free phase by releasing MSG, C/D, I/O, and BSY.

Table 3.1.5-1. Information transfer phases

MSG	Signal		Phase Name	Direction of Transfer	Comment
	C/D	I/O			
0	0	0	DATA OUT	Initiator to target	Data
0	0	1	DATA IN	Initiator from target	Phase
0	1	0	COMMAND	Initiator to target	
0	1	1	STATUS	Initiator from target	
1	0	0	*		
1	0	1	*		
1	1	0	MESSAGE OUT	Initiator to Target	Message
1	1	1	MESSAGE IN	Initiator from Target	Phase

Key: 0 = False, 1 = True, * = Reserved

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 BSY shall remain true and SEL shall remain false. Additionally, during the information transfer phases, the target shall continuously envelope the REQ/ACK handshake(s) with C/D, I/O, and MSG in such a manner that these control signals are valid for a bus settle delay before the assertion of REQ of the first handshake and remain valid until the negation of ACK at the end of the last handshake.

3.1.5.1 Asynchronous information transfer

The target shall control the direction of information transfer by means of the I/O signal. When I/O is true, information shall be transferred from the target to the initiator. When I/O is false, information shall be transferred from the initiator to the target.

If I/O is true (transfer to the initiator), the target shall first drive DB(7-0,P)* to their desired values, delay at least one deskew delay plus a cable skew delay, then assert REQ. DB(7-0,P)* shall remain valid until ACK is true at the target. The initiator shall read DB(7-0,P)* after REQ is true, then signal its acceptance of the data by asserting ACK. When ACK becomes true at the target, the target may change or release DB(7-0, P)* and shall negate REQ. After REQ is false the initiator shall negate ACK. After ACK is false, the target may continue the transfer by driving DB(7-0,P)* and asserting REQ, as described above.

If I/O is false (transfer to the target) the target shall request information by asserting REQ. The initiator shall drive DB(7-0,P)* to their desired values, delay at least one deskew delay plus a cable skew delay and assert ACK. The initiator shall continue to drive the DB(7-0,P)* until REQ is false. When ACK becomes true at the target, the target shall read DB(7-0,P)*, then negate REQ. When REQ becomes false at the initiator, the initiator may change or release DB(7-0,P)* and shall negate ACK. The target may continue the transfer by asserting REQ, as described above.

3.1.5.2 Synchronous data transfer

Synchronous data transfer may be used only in the data phase if previously agreed to by the initiator and target through the message system (see SYNCHRONOUS DATA TRANSFER REQUEST message 3.5.3.2). The messages determine the use of synchronous mode by both SCSI devices and establish a REQ/ACK offset and a transfer period.

The REQ/ACK offset specifies the maximum number of REQ pulses that can be sent by the target 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 by the REQ/ACK offset, the target shall not assert REQ until the next ACK pulse is received. A requirement for successful completion of the data phase is that the number of ACK and REQ pulses be equal.

The target shall assert the REQ signal for a minimum of an assertion period. The target shall wait at least the greater of a transfer period from the last transition of REQ to true or a minimum of a negation period from the last transition of REQ to false before asserting the REQ signal.

The initiator shall send one pulse on the ACK signal for each REQ pulse received. The ACK signal may be asserted as soon as the leading edge of the corresponding REQ pulse has been received. The initiator shall assert the ACK signal for a minimum of an assertion period. The initiator shall wait at least the greater of a transfer period from the last transition of ACK to true or for a minimum of a negation period from the last transition of ACK to false before asserting the ACK signal.

*And where applicable, DB(15-8, P1).

If I/O is asserted (transfer to the initiator), the target shall first drive DB(7-0,P)* to their desired values, wait at least one deskew delay plus one cable skew delay, then assert REQ. DB(7-0,P)* shall be held valid for a minimum of one deskew delay plus one cable skew delay plus one hold time after the assertion of REQ. The target shall assert REQ for a minimum of an assertion period. The target may then negate REQ and change or release DB(7-0,P)*. The initiator shall read the value on DB(7-0,P)* within one hold time of the transition of REQ to true. The initiator shall then respond with an ACK pulse.

If I/O is negated (transfer to the target), the initiator shall transfer one byte for each REQ pulse received. After receiving a REQ pulse, the initiator shall first drive DB(7-0,P)* to their desired values, delay at least one deskew delay plus one cable skew delay, then assert ACK. The initiator shall hold DB(7-0,P)* valid for at least one deskew delay plus one cable skew delay plus one hold time after the assertion of ACK. The initiator shall assert ACK for a minimum of an assertion period. The initiator may then negate ACK and may change or release DB(7-0,P)*. The target shall read the value of DB(7-0,P)* within one hold time of the transition of ACK to true.

3.1.6 Command phase

The Command phase allows the target to request command information from the initiator.

The target shall assert the C/D signal and negate the I/O and MSG signals during the REQ/ACK handshake(s) of this phase.

3.1.7 Data phase

The Data phase is a term that encompasses both the Data In phase and the Data Out phase.

3.1.7.1 Data in phase

The Data In phase allows the target to request that it send data to the initiator.

The target shall assert the I/O signal and negate the C/D and MSG signals during the REQ/ACK handshake(s) of this phase.

3.1.7.2 Data out phase

The Data Out phase allows the target to request that data be sent to it from the initiator.

The target shall negate the C/D, I/O, and MSG signals during the REQ/ACK handshake(s) of this phase.

3.1.8 Status phase

The Status phase allows the target to request that it send status information to the initiator.

See section 4.3 for details.

The target shall assert C/D and I/O and negate the MSG signal during the REQ/ACK handshake of this phase.

*And where applicable, DB(15-8, P1).

3.1.9 Message phase

The Message phase is a term that references either a Message In or a Message Out phase. Multiple messages may be sent during either phase. Multiple byte messages shall be wholly contained with a single message phase. Messages supported by a particular drive are listed in the Product Manual for that drive.

3.1.9.1 Message in phase

The Message In phase allows the target to request that it send message(s) to the initiator.

The target shall assert C/D, I/O, and MSG during the REQ/ACK handshake(s) of this phase.

3.1.9.2 Message out phase

The Message Out phase allows the target to request that message(s) be sent from the initiator to the target. The target may invoke this phase at its convenience in response to the Attention condition (see 3.2.1) created by the initiator.

The target shall assert C/D and MSG and negate I/O during the REQ/ACK handshake(s) of this phase. The target shall handshake byte(s) in this phase until ATN goes false, unless an error occurs (see Message Reject, 3.5.2).

If the target detects one or more parity error(s) on the message byte(s) received, it may indicate its desire to retry the message(s) by asserting REQ after detecting ATN has gone false and before changing to any other phase. The initiator, upon detecting this condition, shall resend all of the previous message byte(s) sent during this phase. When resending more than one message byte, the initiator shall assert ATN before asserting ACK on the first byte and shall maintain ATN asserted until the last byte is sent as described in 3.2.1.

If the target receives all of the message byte(s) successfully (i.e., no parity errors), it shall indicate that it shall not retry by changing to any information transfer phase other than the Message Out phase and transfer at least one byte. The target may also indicate that it has successfully received the message byte(s) by changing to the Bus Free phase (e.g., Abort or Bus Device Reset messages).

3.1.10 Signal restrictions between phases

When the SCSI bus is between two information transfer phases, the following restrictions shall apply to the SCSI bus signals:

1. The BSY, SEL, REQ, and ACK signals shall not change.
2. The C/D, I/O, MSG, and Data Bus signals may change. When switching the Data Bus direction from Out (initiator driving) to In (target driving), the target shall delay driving the Data Bus by at least a data release delay plus settle delay after asserting the I/O signal and the initiator shall release 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 direction from In (target driving) to Out (initiator driving), the target shall release the Data Bus no later than a deskew delay after negating the I/O signal.
3. The ATN and RST signals may change as defined under the descriptions for the Attention condition (3.2.1) and Reset condition (3.2.2).

3.2 SCSI bus conditions

The SCSI bus has three asynchronous conditions; the Attention condition, the Reset condition and the Contingent Allegiance condition. These conditions cause the SCSI device to perform certain actions and can alter the phase sequence.

3.2.1 Attention condition

The Attention condition allows an initiator to inform a target that the initiator has a message ready. The target gets this message at its convenience by performing a Message Out phase.

The initiator creates the Attention condition by asserting ATN at any time except during the Arbitration or Bus Free.

The initiator shall assert the ATN signal two deskew delays before negating the ACK signal for the last byte transferred in a bus phase for the attention condition to be honored before transition to a new bus phase. Asserting the ATN signal later might not be honored until a later bus phase and then may not result in the expected action. The initiator shall negate the ATN signal two deskew delays before asserting the ACK signal while transferring the last byte of the message. If the target detects that the initiator failed to meet this requirement, then the target shall go to BUS FREE phase (see unexpected BUS FREE, 3.1.1).

The drive responds with MESSAGE OUT phase as follows:

1. If ATN occurs during a Data phase, Message Out occurs at a convenient time. It may not occur until several logical blocks after ATN is first asserted.
2. If ATN occurs during a Command phase, Message Out occurs after transfer of all Command Descriptor Block bytes has been completed.
3. If ATN occurs during a Status phase, Message Out occurs after the status byte has been acknowledged by the initiator.
4. If ATN occurs during a Message In Phase, Message Out occurs after the last byte of the current message has been acknowledged by the initiator.
5. If ATN occurs during a Selection or Reselection phase, Message Out occurs immediately after that Selection or Reselection phase. In the Reselection case, the drive enters the Message Out phase after it has sent its Identify message for that Reselection Phase.

The initiator shall keep ATN asserted if more than one byte is to be transferred. The initiator may negate the ATN signal at any time except while the ACK signal is asserted during a Message Out phase. Recommended practice is that the initiator negates ATN while REQ is true and ACK is false during the last REQ/ACK handshake of the Message Out phase.

3.2.2 Reset condition

The Reset condition is used to immediately clear all SCSI devices from the bus. This condition shall take precedence over all other phases and conditions. During the Reset condition, the state of all SCSI bus signals other than RST is not defined.

The drive never asserts the Reset signal.

All SCSI devices shall release all SCSI bus signals (except RST) within a bus clear delay of the transition of RST to true. The Bus Free phase always follows the Reset condition.

Drives that operate in accordance with this specification implement only the “hard” Reset option. Upon detection of the Reset condition, the drive shall:

1. Clear all uncompleted commands,
2. Release all SCSI device reservations,
3. Return any SCSI device operating modes (Mode Select, etc.) to their default or last saved conditions.
4. Activate Unit Attention Condition for all Initiators.

3.2.3 Contingent allegiance condition

The Contingent Allegiance condition shall exist following the return of CHECK CONDITION or COMMAND TERMINATED STATUS and may optionally exist following an unexpected disconnect. The contingent allegiance condition shall be preserved for the I T x nexus until it is cleared. The contingent allegiance condition shall be cleared upon the generation of a hard reset condition or by an ABORT message, a BUS DEVICE RESET message, or any subsequent Untagged command for the I T x nexus. While the contingent allegiance condition exists the drive shall preserve the sense data for the initiator.

While the contingent allegiance condition exists, the drive shall respond to any other requests for access to the logical unit from another initiator with a BUSY status. Execution of queued commands shall be suspended until the Contingent Allegiance condition is cleared.

While the Contingent Allegiance exists, if the initiator illegally issues a tagged command, Seagate drives shall return a BUSY status to the initiator and maintain the Contingent Allegiance.

See individual drive Product manual to see if a particular drive supports the Contingent Allegiance condition.

3.3 SCSI bus phase sequences

The order in which phases are used on the SCSI bus follows a prescribed sequence.

In all systems, the Reset condition can abort any phase and is always followed by the Bus Free phase. Also, any other phase can be followed by the Bus Free phase.

3.3.1 Nonarbitrating system

For systems in which the Arbitration phase is not implemented, the allowable sequences are shown in Figure 3.3-2. The normal progression is from the Bus Free phase to Selection, and from Selection to one or more of the information transfer phases (Command, Data, Status, or Message).

3.3.2 Arbitrating systems

For systems in which the Arbitration phase is implemented, the allowable sequences are shown in Figure 3.3-1. The normal progression is from the Bus Free phase to Arbitration, from Arbitration to Selection or Reselection, and from Selection or Reselection to one or more of the information transfer phases (Command, Data, Status, or Message).

3.3.3 All systems

There are no restrictions on the sequences between information transfer phases. A phase type may even be followed by the same phase type (e.g., a Data phase may be followed by another Data phase).

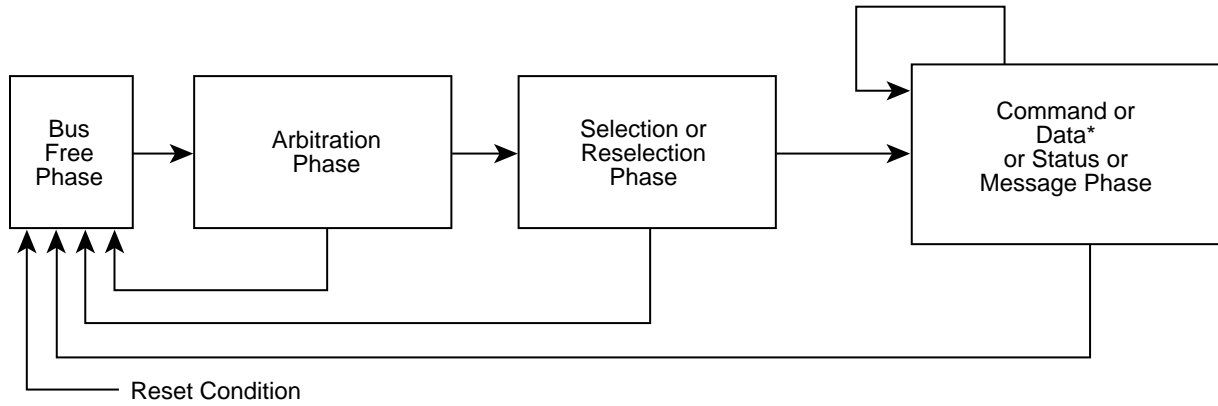


Figure 3.3-1. Phase sequences with arbitration

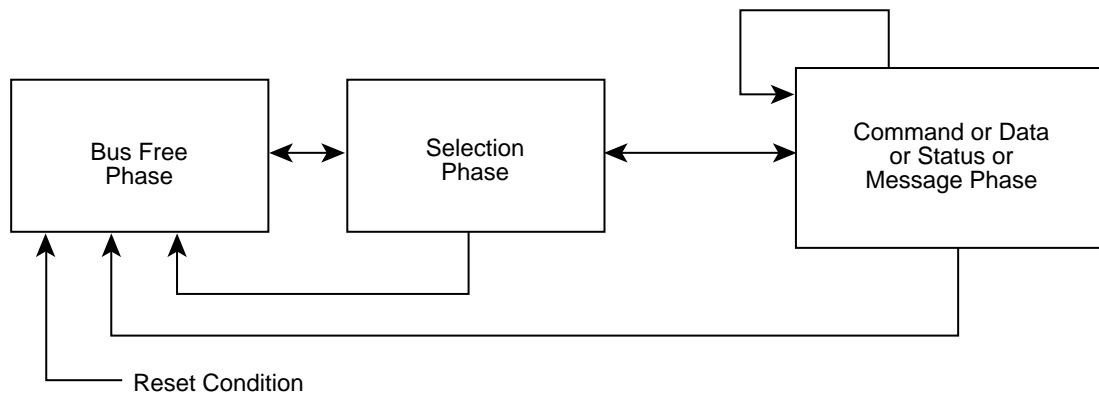


Figure 3.3-2. Phase sequences without arbitration

*Data phase must follow a command phase except after reselection phase

3.4 SCSI pointers

The drive supports systems that use the pointer philosophy described in the following paragraphs.

Consider the system shown in Figure 3.4-1 in which an initiator and target communicate on the SCSI bus in order to execute a command.




Figure 3.4-1. Simplified SCSI system

The SCSI architecture provides for sets of pointers to be kept within each initiator Path Control area (see Figure 3.4-1). These pointers are in sets of three pointers per set. The pointers in each set point to three storage area sections in the initiator. The three sections contain the following information:

1. A command from initiator to target.
2. Status (from target) associated with the command.
3. Data (to/from target) associated with the command.

Of these three-pointer sets there are two types:

1. Current (active) pointers (one set per initiator only).
2. Saved pointers (one or more sets per initiator, up to seven sets total).

The use of these two types of pointers is described in the following paragraphs.

3.4.1 Current pointers

Current Pointers represent the current state of the interface between the initiator and the target the initiator is currently connected to and servicing. The pointers for the current command remain in the initiator Current Pointer registers from the time they are put there after the completion of all activities associated with the previous command until the logic of the Initiator dictates a new command shall be executed. Normally, successful receipt by the initiator of good status associated with the current command triggers the initiator to insert a new set of the three “current” pointers for the next command. The initiator does not wait for the Command Complete signal before deciding whether to retry the current command or transfer in new command pointers. If the current command was never satisfactorily completed, the initiator logic may dictate that some special response action be taken, such as restoring the values in the current pointer registers to their beginning values so the current command can be resent*, or sending a command such as a Request Sense command to the target, or the initiator could ignore the unsatisfactorily completed command and send the next originally scheduled command.

* For example, if the drive detects a parity error in the data out from the current command, it sends the “Restore Pointers” message to the initiator. In this case, the Restore Pointers request causes the initiator to restore the current pointers to the values existing at the beginning of the current command so the current command can be resent. The “beginning” pointer values point to the first byte of the current Command Descriptor Block, the first byte of the area set aside for status to be returned and the first byte of the area set aside for data associated with the current command. (See paragraph 3.4.2 for a detailed description of operations resulting from the Restore Pointers message).

3.4.2 Saved pointers

Saved pointers point to initiator storage locations where command, status and data information are stored for a command that was saved at some point in the past. There is one set of saved pointers for the current command for each target on the interface bus that is currently active (whether or not it is currently connected to the initiator). The saved command pointer always points to the first byte of the Command Descriptor Block (see paragraph 4.2) for the “current” command for each target, the saved status pointer always points to the first byte of the area used for the status associated with the command, and the saved data pointer points to some location (not necessarily the beginning) in the area used for data associated with the command. When a target disconnects the initiator saves the current pointers. Before a target disconnects it may send a Save Data Pointers message to the initiator, which copies the data pointer that is for the current command for that target into the location set aside for the target’s saved pointers. When a target reconnects, the initiator performs a restore pointers operation that copies the saved pointers for the reconnected target into the initiator current pointer registers so that the current command for that target may continue its operation from where it left off before disconnecting. If the target had sent a Save Data Pointer message previously, the current data pointer points to the place in the data store area where operations left off. The data pointer otherwise points to the beginning of the data area, unless the data pointers were modified by a MODIFY DATA POINTERS message from the target prior to disconnecting. The MODIFY DATA POINTERS message adds a value to the data pointer that allows data to be taken, upon reconnection, from a location before or after the last byte transferred location.

3.5 Message system specification

The message system allows communication between an initiator and target for the purpose of interface management. For the drive, the Logical Unit Number (see Section 4.2.2), (LUN) is always zero.

3.5.1 General message protocol

A message may be one, two, or multiple bytes in length. One or more messages may be sent during a single MESSAGE phase, but a message may not be split over MESSAGE phases. The initiator is required to end the MESSAGE OUT phase (by negating ATN) when it sends certain messages identified in Table 3.5.2-1. One-byte, Two-byte, and extended message formats are defined. The first byte of the message determines the format as follows:

Value	Message Format
00h	One-Byte Message (Command Complete)
01h	Extended Messages
02h - 1Fh	One-Byte Messages
20h - 2Fh	Two-Byte Messages
30h - 7Fh	Reserved
80h - FFh	One-Byte Message (Identify)

One-Byte messages consist of a single byte transferred during a MESSAGE phase. The value of the byte determines the message to be performed as defined in Table 3.5.2-1.

Two-byte messages consist of two consecutive bytes transferred during a MESSAGE phase. The value of the first byte determines the message to be performed as defined in Table 3.5.2-1. The second byte is a parameter byte that is used as defined in the message descriptions in Section 3.5.2.

See section 3.5.3.1 for details on the one and two byte messages.

Extended messages consist of from three to 7 bytes transferred during the MESSAGE phase. The extended message includes a three byte header and up to 4 extended message argument bytes. See Section 3.5.3.2 for details of the extended messages.

3.5.2 Messages - General

The messages supported by the various drives are listed in a table in the drive's Product Manual. Table 3.5.2-1 lists Messages that are defined by the SCSI protocol. The message code values are given a direction specification (in-Out). Detailed descriptions follow the table. Messages other than those supported by a drive are answered by the drive with a Message Reject message.

Table 3.5.2-1. Message codes

Code	Message Name	Direction		Negate Atn Before
		In	Out	Last Ack
01h***	EXTENDED MESSAGES	In	Out	Yes
06h	ABORT		Out	Yes
0Dh	ABORT TAG		Out	Yes
0Ch	BUS DEVICE RESET		Out	Yes
0Eh	CLEAR QUEUE		Out	Yes
00h	COMMAND COMPLETE	In		—
04h	DISCONNECT	In		—
80h - FFh	IDENTIFY	In		—
80h - FFh	IDENTIFY		Out	No
23h	IGNORE WIDE RESIDUE (Two Bytes)	In		—
0Fh	INITIATE RECOVERY	In	Out	Yes
05h	INITIATOR DETECTED ERROR		Out	Yes
0Ah	LINKED COMMAND COMPLETE	In		—
0Bh	LINKED COMMAND COMPLETE (With Flag)	In		—
09h	MESSAGE PARITY ERROR		Out	Yes
07h	MESSAGE REJECT [1]	In	Out	Yes
***	MODIFY DATA POINTER	In		—
08h	NO OPERATION		Out	Yes
	Queue Tag Messages (Two Bytes)			
21h	HEAD OF QUEUE TAG		Out	No
22h	ORDERED QUEUE TAG		Out	No
20h	SIMPLE QUEUE TAG	In	Out	No
10h	RELEASE RECOVERY		Out	Yes
03h	RESTORE POINTERS	In		—
02h	SAVE DATA POINTER	In		—
***	SYNCHRONOUS DATA TRANSFER REQUEST			
		In	Out	Yes
***	WIDE DATA TRANSFER REQUEST			
		In	Out	Yes
11h	TERMINATE I/O PROCESS		Out	Yes
12h	CONTINUE I/O PROCESS		Out	Yes
13h	TARGET TRANSFER DISABLE		Out	Yes
14h - 1Fh	RESERVED			
24h - 2Fh	RESERVED for two-byte messages			
30h - 7Fh	RESERVED			

KEY: In = Target to initiator. Out = Initiator to target
 Yes = Initiator shall negate ATN before last ACK of message
 No = Initiator may or may not negate ACK before last ACK of message. (see attention condition)
 — = Not Applicable
 *** = Extended message (see Tables 3.5.3-3 through 3.5.3-5)
 [1] = The drive does not resend (RETRY) the original message.

The first message sent by the initiator after the SELECTION phase shall be an IDENTIFY, ABORT, or BUS DEVICE RESET message. If a target receives any other message it shall go to BUS FREE phase (unexpected BUS FREE).

If the first message is an IDENTIFY message, then it may be immediately followed by other messages, such as the first of a pair of SYNCHRONOUS DATA TRANSFER REQUEST messages. If tagged queuing is used the queue tag message immediately follows the IDENTIFY message. The IDENTIFY message establishes a logical connection between the initiator and the specified logical unit known as an I T L nexus.

After the RESELECTION phase, the target's first message shall be IDENTIFY. This allows the I T L nexus to be reestablished. Only one logical unit shall be identified for any connection; if a target receives a second IDENTIFY message with a logical unit number other than zero during a connection, it shall go to BUS FREE phase (unexpected BUS FREE).

Whenever an I T L nexus is established by an initiator that is allowing disconnection, the initiator shall ensure that the active pointers are equal to the saved pointers for that particular logical unit. An implied restore pointers operation shall occur as a result of a reconnection.

3.5.3 Message details

SCSI message details are defined in the following paragraphs.

3.5.3.1 One and two byte messages

Abort (06h)

The *Abort* message is sent from the initiator to the target to clear the present I/O process plus any queued I/O process for the I T L nexus. The target shall go to the BUS FREE phase following successful receipt of this message. Pending data, status, and queued I/O processes for any other I T L nexus shall not be cleared.

If only an I T nexus has been established, the target shall go to the BUS FREE phase. No status or message shall be sent for the I/O process and the I/O process queue shall not be affected.

It is not an error to issue this message to an I T L nexus that does not currently have an active or queued I/O process. Transmission of this message shall terminate any extended contingent allegiance condition that may exist between the I T L nexus.

Notes.

1. The BUS DEVICE RESET, CLEAR QUEUE, ABORT, and ABORT TAG messages provide a means to clear one or more I/O processes prior to normal termination. The BUS DEVICE RESET message clears all I/O processes for all initiators on all logical units of the target. The CLEAR QUEUE message clears all I/O processes for all initiators on the specified logical unit of the target. The ABORT message clears all I/O processes for the selecting initiator on the specified logical unit of the target. The ABORT TAG message clears the current I/O process only.
2. It is permissible for an initiator to select a target that is currently disconnected for the purpose of sending one of the above message sequences.

Abort Tag (0Dh)

The ABORT TAG message shall be implemented if tagged queuing is implemented. The target shall go to the BUS FREE phase following successful receipt of this message. The target shall clear the current I/O process for the I T L nexus. If the target has already started execution of the I/O process, the execution shall be halted. The medium contents may have been modified before the execution was halted. In either case, any pending status or data for the I/O process shall be cleared and no status or ending message shall be sent to the initiator. Pending status, data, and commands for other queued or executing I/O processes shall not be affected. Execution of other I/O processes queued for the I T L nexus shall continue in the normal manner.

Bus Device Reset (0Ch)

The BUS DEVICE RESET message is sent from an initiator to direct a target to clear all current I/O processes on that SCSI device. This message forces a hard reset condition to the selected SCSI drive. The target shall go to the BUS FREE phase following successful receipt of this message. The target shall create a Unit Attention condition for all initiators (See 4.6).

Clear Queue (0Eh)

The CLEAR QUEUE message shall be implemented if tagged queuing is implemented and may be implemented if untagged queuing is implemented. The target shall go to the BUS FREE phase following successful receipt of this message. The target shall perform an action equivalent to receiving a series of ABORT messages from each initiator. All I/O processes, from all initiators, in the queue for the specified logical unit shall be cleared from the queue. All executing I/O processes shall be halted. The medium may have been altered by partially executed commands. All pending status and data for that logical unit for all initiators shall be cleared. No status or ending message shall be sent for any of the outstanding I/O processes. A unit attention condition shall be generated for all other initiators with I/O processes that either had been executing or were queued for execution for that logical unit. When reporting the Unit Attention condition the additional sense code shall be set to TAGGED COMMANDS CLEARED BY ANOTHER INITIATOR.

Command Complete (00h)

The COMMAND COMPLETE message is sent from a target to an initiator to indicate that the execution of a command (or series of linked commands) has terminated and that valid status has been sent to the initiator. After successfully sending this message, the target shall go to the BUS FREE phase by releasing BSY. The target shall consider the message transmission to be successful when it detects the negation of ACK for the COMMAND COMPLETE message with the ATN signal false.

Notes: The command may have been executed successfully or unsuccessfully as indicated in the status.

Continue I/O Process (12h)

The CONTINUE I/O PROCESS message is sent from the initiator to the target to reconnect to an I/O process. This message shall be sent in the same MESSAGE OUT phase as the IDENTIFY message.

IMPLEMENTORS NOTE: Thus the MESSAGE OUT phase following SELECTION phase consists of the IDENTIFY, queue tag (if any), and CONTINUE I/O PROCESS messages.

The purpose of the CONTINUE I/O PROCESS message is to distinguish a valid initiator reconnection from an incorrect initiator reconnection (see 4.9.2).

If the target expects a significant delay before it will be ready to continue processing the reconnected I/O PROCESS, it may attempt to free the SCSI bus by sending a DISCONNECT message to the initiator. The initiator may reject the disconnection attempt by responding with MESSAGE REJECT message.

It is an error for the initiator to send this message on an initial connection (i.e., there is no I/O process for the nexus) and the target shall go to the BUS FREE phase (see unexpected disconnect, 3.1.1).

Initiators should avoid sending this message to targets which have not implemented this message. Such targets may not respond as described in this section. An initiator can determine whether a target implements this message by examining the TranDis bit in the standard INQUIRY data (see Table 5.1.1-8).

Disconnect (04h)

The DISCONNECT message is sent from a target to inform an initiator that the present connection is going to be broken (the target plans to disconnect by releasing the BSY signal), but that a later reconnect will be required in order to complete the current I/O process. This message shall not cause the initiator to save the data pointer. After successfully sending this message, the target shall go to the BUS FREE phase by releasing the BSY signal. The target shall consider the message transmission to be successful when it detects the negation of the ACK signal for the DISCONNECT message with the ATN signal false.

Targets which break data transfers into multiple connections shall end each successful connection (except possibly the last) with a SAVE DATA POINTER - DISCONNECT message sequence.

Identify (80h - FFh)

The IDENTIFY message (Table 3.5.3-1) is sent by either the initiator or the target to establish an I T L nexus.

Table 3.5.3-1. Identify message format

Bit Byte	7	6	5	4	3	2	1	0
0	Identify	DiscPriv	LUN- TAR	LUNTRN				

The identify bit shall be set to one to specify that this is an IDENTIFY message.

A disconnect privilege (DiscPriv) bit of one specifies that the initiator has granted the target the privilege of disconnecting. A DiscPriv bit of zero specifies that the target shall not disconnect. Those drive models that do not implement this feature must set this bit to zero. This bit is used in connection with the Disconnect Immediate (DImm) bit in the Disconnect Reconnect Mode Page.

Note.

If an initiator selects a drive using the SCSI-1 Single Initiator Option, the drive shall not attempt to disconnect from that I/O process, regardless of the state of the "Disconnect Privilege" bit in the Identify message (if sent).

A logical unit target (LUNTAR) bit of zero specifies that the IDENTIFY message is directed to a logical unit. A LUNTAR bit of one specifies that the IDENTIFY message is directed to a target routine that does not involve the logical unit. This bit is not supported by the drive and will always be interpreted as a 0.

The logical unit number target routine number (LUNTRN) field specifies a logical unit number if the LUNTAR bit is zero. The LUNTRN field specifies a target routine number if the LUNTAR bit is one. Only the INQUIRY and REQUEST SENSE commands are valid for target routines. If a target receives any other command for a target routine, it shall return CHECK CONDITION status and shall set the sense key to ILLEGAL REQUEST.

An invalid LUNTRN field terminates the command with a CHECK CONDITION, and in response to a REQUEST SENSE command the drive returns sense data with the sense key set to ILLEGAL REQUEST and the additional sense code set to LOGICAL UNIT NOT SUPPORTED.

An IDENTIFY message is invalid if a reserved bit is set to one or if the LUNTAR bit is set to one and the target does not implement target routines. A device may respond to an invalid IDENTIFY message by immediately sending a MESSAGE REJECT message or by returning CHECK CONDITION status. If a CHECK CONDITION status is returned, the sense key shall be set to ILLEGAL REQUEST and the additional sense code shall be set to INVALID BITS IN IDENTIFY MESSAGE FIELD.

If a valid IDENTIFY message has not yet been received for the current nexus, the drive returns a MESSAGE REJECT message and goes to the BUS FREE phase (see 3.1.1, unexpected disconnect).

Only one logical unit number or target routine number shall be identified per I/O process. The initiator may send one or more IDENTIFY messages during a connection. A second IDENTIFY message with a different value in either the LUNTAR bit or LUNTRN field shall not be issued before a BUS FREE phase has occurred; if a target receives a second IDENTIFY message with a different value in either of these fields, it shall go to BUS FREE phase (unexpected BUS switch to another I/O process. (See the DTDC field of the disconnect-reconnect page (5.2.1-23) for additional controls over disconnection.)

An implied RESTORE POINTERS message shall be performed by the initiator prior to the assertion of the ACK signal on the next phase for an IDENTIFY message sent during reconnection.

Initiator Detected Error (05h)

The INITIATOR DETECTED ERROR message is sent from an initiator to inform a target that an error has occurred that does not preclude the target from retrying the operation. The source of the error may either be related to previous activities on the SCSI bus or may be internal to the initiator and unrelated to any previous SCSI bus activity. Although present pointer integrity is not assured, a RESTORE POINTERS message or a disconnect followed by a reconnect, shall cause the pointers to be restored to their defined prior state.

Linked Command Complete (0Ah)

The LINKED COMMAND COMPLETE message is sent from a target to an initiator to indicate that the execution of a linked command has completed and that status has been sent. The initiator shall then set the pointers to the initial state for the next linked command.

Linked Command Complete (with Flag) (0Bh)

The LINKED COMMAND COMPLETE (WITH FLAG) message is sent from a target to an initiator to indicate that the execution of a linked command (with the flag bit set to one) has completed and that status has been sent. The initiator shall then set the pointers to the initial state of the next linked command.

Message Parity Error (09h)

The MESSAGE PARITY ERROR message is sent from the initiator to the target to indicate that the last message byte it received had a parity error.

In order to indicate its intentions of sending this message, the initiator shall assert the ATN signal prior to its release of the ACK signal for the REQ/ACK handshake of the message that has the parity error. This provides an interlock so that the target can determine which message has the parity error. If the target receives this message under any other circumstance, it shall signal a catastrophic error condition by releasing the BSY signal without any further information transfer attempt (see 3.1.1).

Message Reject (07h)

The MESSAGE REJECT message is sent from either the initiator or target to indicate that the last message byte it received was inappropriate or has not been implemented.

In order to indicate its intentions of sending this message, the initiator shall assert the ATN signal prior to its release of the ACK signal for the REQ/ACK handshake of the message byte that is to be rejected. If the target receives this message under any other circumstance, it shall reject this message.

When a target sends this message, it shall change to MESSAGE IN phase and send this message prior to requesting additional message bytes from the initiator. This provides an interlock so that the initiator can determine which message byte is rejected.

Note. After a target sends a MESSAGE REJECT message and if the ATN signal is still asserted, then it returns to the MESSAGE OUT phase. The subsequent MESSAGE OUT phase begins with the first byte of a message, not the middle of the previous message.

Modify Data Pointer (01h)

See Paragraph 3.5.3.2 on Extended Messages

No Operation (08h)

The NO OPERATION message is sent from an initiator in response to a target's request for a message when the initiator does not currently have any other valid message to send.

For example, if the target does not respond to the attention condition until a later phase and at that time the original message is no longer valid the initiator may send the NO OPERATION message when the target enters the MESSAGE OUT phase.

Queue Tag Messages (20h, 21h or 22h)

When one or more initiators have multiple I/O processes to be queued by a target, each I/O process must have a queue tag associated with it. The queue tag is specified in a Queue Tag Message that must accompany the initiation of each such I/O process. See also Section 4.7.2, “Tagged Queuing”.

Table 3.5.3-2. Queue tag message format

Bit Byte	7	6	5	4	3	2	1	0
0	Message Code (20h or 21h or 22h)							
1	Queue Tag							

Table 3.5.3-2 defines the format for the queue tag messages. Those drives that implement tagged queuing and use the queue tag messages HEAD OF QUEUE TAG, ORDERED QUEUE TAG, and SIMPLE QUEUE TAG indicate so in their individual Product Manuals.

The queue tag messages are used to specify an identifier, called a queue tag, for an I/O process that establishes the I L Q nexus. The queue tag field is an 8-bit unsigned integer assigned by the initiator during an initial connection. The queue tag for every I/O process for each I_T L nexus should be unique. If the drive receives a queue tag identical to one that is currently in use for the I T L nexus, the drive shall abort ALL I/O processes for the initiator and shall return CHECK CONDITION status. The sense key shall be set to ABORTED COMMAND and the additional sense code shall be set to OVERLAPPED COMMANDS ATTEMPTED. Only one status is returned.

Note. For each logical unit on each target, each initiator has up to 256 queue tags to assign to I/O processes. A queue tag becomes available for reassignment when the I/O process ends. Typically drives to which this specification applies queue up to 64 tagged I/O processes concurrently for up to seven initiators on the bus. There could be 64 queue tags for one initiator, or the 64 can be divided among the seven possible initiators, not necessarily evenly. See applicable drive Product Manual for number of queue tags supported.

The appropriate queue tag message shall be sent immediately following the IDENTIFY message and within the same message phase to establish the I T L Q nexus for the I/O process. Only one I T L Q nexus may be established during a connection. If a queue tag message is not sent, then only an I T L nexus is established for the I/O process (untagged command).

If a target attempts to reconnect using an invalid queue tag, then the initiator should respond with an ABORT TAG message.

- **Head of Queue Tag (21h)**

The *Head of Queue Tag* message specifies that the I/O process be placed first in that logical unit's queue for execution. An I/O process already being executed by the target shall not be preempted. A subsequent I/O process received with a HEAD OF QUEUE TAG message shall be placed at the head of the queue for execution in last-in, first-out order.

- **Ordered Queue Tag (22h)**

The ORDERED QUEUE TAG message specifies that the I/O process be placed in the drive's I/O process queue for execution in the order received, with respect to other commands with ORDERED QUEUE TAG messages, except for I/O processes received with a HEAD OF QUEUE TAG message, which are placed at the head of the queue.

- **Simple Queue Tag (20h)**

The SIMPLE QUEUE TAG message specifies that the I/O process be placed in the drive's I/O process queue for execution. The order of execution can be arranged by the drive in accordance with a performance optimization algorithm. The SIMPLE QUEUE TAG message is also sent by the target when reconnecting to the initiator.

Restore Pointers (03h)

The RESTORE POINTERS message is sent from a target to direct the initiator to restore the most recently saved pointers (for the current nexus) to the active state. Pointers to the command, data, and status locations for the nexus shall be restored to the active pointers. Command and status pointers shall be restored to the beginning of the present command and status areas. The data pointer shall be restored to the value at the beginning of the data area in the absence of a SAVE DATA POINTER message or to the value at the point where the last SAVE DATA POINTER message occurred for that nexus.

Save Data Pointer (02h)

The SAVE DATA POINTER message is sent from a target to direct the initiator to save a copy of the present active data pointer for the current nexus. (See 3.4 for a definition of pointers.)

Synchronous Data Transfer (01h)

See Paragraph 3.5.3.2 on EXTENDED MESSAGE.

Terminate I/O Process (11h)

This message is supported only on those drives having it as a factory installed option.

The Terminate I/O Process message is sent from the initiator to the target to terminate the current I/O process without corrupting the medium.

With the following exceptions, the target shall terminate the current I/O process and return COMMAND TERMINATED status. The sense key shall be set to NO SENSE. The additional sense code and qualifier are set to I/O PROCESS TERMINATED.

If the associated I/O process involves a data phase, the target shall set the valid bit in the sense data to one and set the information field as follows:

- 1) If the command descriptor block specifies an allocation length or parameter list length, the information field shall be set to the difference (residue) between the number of bytes successfully transferred and the requested length.
- 2) If the command descriptor block specifies a transfer length field, the information field shall be set as defined in the REQUEST SENSE command (see 5.1.1.2).

If an error is detected for the associated I/O process the target shall ignore the TERMINATE I/O PROCESS message.

If the operation requested for the associated I/O process has been completed but status has not been returned, the target shall ignore the TERMINATE I/O PROCESS message.

If the target does not support this message or is unable to stop the current I/O process, it shall send a MESSAGE REJECT message to the initiator and continue the I/O process in a normal manner.

The effect of a TERMINATE I/O PROCESS message on the command queue depends on the queue error recovery option specified in the control mode page (see Table 5.2.1-28) and on whether or not a contingent allegiance condition is generated.

Target Transfer Disable (13h)

The TARGET TRANSFER DISABLE (TDD) message is sent from an initiator to a target to request that subsequent reconnections for data transfer on the I/O process be done by the initiator instead of the target. The target may reconnect for other purposes, but shall not enter a data phase on a target reconnection. SCSI devices that implement this message shall also implement the CONTINUE I/O PROCESS message.

If used, this message shall be sent as the last message of the first MESSAGE OUT phase of an initial connection. The target may continue the I/O process, including any DATA OUT phases on the initial connection, until the target would normally disconnect, but the target shall not reconnect to transfer data. That is, the target shall not enter a DATA IN phase on the initial connection and the target shall not enter any data phase on any subsequent target reconnection for the I/O process.

When the target is ready to transfer data for a disconnected I/O process for which a TTD message has been sent, the target shall reconnect to the initiator for the I/O process (via a RESELECTION phase, an IDENTIFY message, and an optional SIMPLE QUEUE TAG message), send a DISCONNECT message, and, if the initiator does not respond with a MESSAGE REJECT message, go to the BUS FREE phase. This connection serves to notify the initiator that the I/O process is ready for data transfer. If the initiator rejects the DISCONNECT message, the target may enter a data phase; otherwise, the initiator may reconnect to the I/O process as described in the CONTINUE I/O PROCESS message (see above) to do the data transfer.

Initiators should avoid sending the TTD message to targets which have not implemented this message. Such targets may not respond as described in this section. An initiator can determine whether a target implements this message by examining the TranDis bit in the standard INQUIRY data (see 5.1.1.3).

3.5.3.2 Extended message (01h)

Extended messages are messages that require more than two bytes to send the necessary information. The Extended messages supported by the drive are Modify Data Pointer and Synchronous Data Transfer. These two are described in detail in the following paragraphs.

MODIFY DATA POINTERS MESSAGE

Table 3.5.3-3. Modify data pointer

Byte	Value	Description
0	01h	Extended message
1	05h	Extended message length
2	00h	MODIFY DATA POINTER code
3		Argument (Most Significant Byte)
4		Argument
5		Argument
6		Argument (Least Significant Byte)

The MODIFY DATA POINTER message (Table 3.5.3-3) is sent from the target to the initiator and requests that the signed argument be added (two's complement) to the value of the current data pointer.

SYNCHRONOUS DATA TRANSFER REQUEST MESSAGE

Table 3.5.3-4. Synchronous data transfer request

Byte	Value	Description
0	01h	Extended message
1	03h	Extended message length
2	01h	SYNCHRONOUS DATA TRANSFER REQUEST code
3	m	Transfer period (m times 4 nanoseconds) [1]
4	x	REQ/ACK offset

[1] See Tables in the individual drive Product Manuals, interface section, for a list of transfer periods supported by the drive described therein.

A SYNCHRONOUS DATA TRANSFER REQUEST (SDTR) message (Table 3.5.3-4) exchange shall be initiated by a SCSI device whenever a previously arranged data transfer agreement may have become invalid. The agreement becomes invalid after any condition that may leave the data transfer agreement in an indeterminate state such as

1. after a hard reset condition
2. after a BUS DEVICE RESET message and
3. after a power cycle.
4. for SCSI devices that implement WDTR, occurrence of an intervening WDTR.

In addition, a SCSI device may initiate an SDTR message exchange whenever it is appropriate to negotiate a new data transfer agreement (either synchronous or asynchronous). SCSI devices that are capable of synchronous data transfers shall not respond to an SDTR message with a MESSAGE REJECT message.

The SDTR message exchange establishes the permissible transfer periods and REQ/ACK offsets for all logical units on the two devices.

The transfer period is the minimum time allowed between leading edges of successive REQ pulses and of successive ACK pulses to meet the device requirements for successful reception of data.

The 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 at the target. This value is chosen to prevent overflow conditions in the device's reception buffer and offset counter. A REQ/ACK offset value of zero shall indicate asynchronous data transfer mode; a value of FFh shall indicate unlimited REQ/ACK offset. The REQ/ACK offset value supported by each drive is listed in each drive product manual, Volume 1.

The originating device (the device that sends the first of the pair of SDTR messages) sets its values according to the rules above to permit it to receive data successfully. If the responding device can also receive data successfully with these values, it returns the same values in its SDTR message. If it requires a larger transfer period, smaller REQ/ACK offset, or both in order to receive data successfully, it substitutes values in its SDTR message as required, returning unchanged any value not required to be changed. Each device when transmitting data shall respect the limits set by the other's SDTR message but it is permitted to transfer data with larger transfer periods, small REQ/ACK offsets, or both than specified in the other's SDTR message. The successful completion of an exchange of SDTR messages implies an agreement as follows:

Responding Device SDTR SDTR response

Implied Agreement

- | | |
|---|---|
| 1. Non-zero REQ/ACK offset | Each device transmits data with a transfer period equal to or less than the values received in the other device's SDTR message. |
| 2. REQ/ACK offset equal to zero | Asynchronous transfer |
| 3. MESSAGE REJECT MESSAGE | Asynchronous transfer |
| 4. No Response from Initiator to drive Initiated SDTR | Asynchronous transfer |

If the initiator recognizes that negotiation is required, it asserts the ATN signal and sends an SDTR message to begin the negotiating process. After successfully completing the MESSAGE OUT phase, the target shall respond with the proper SDTR message. If an abnormal condition prevents the target from returning an appropriate response, both devices shall go to asynchronous data transfer mode for data transfers between the two devices.

Following target response (1) above, the implied agreement for synchronous operation shall be considered to be negated by both the initiator and the target if the initiator asserts the ATN signal and the first message out is either MESSAGE PARITY ERROR or MESSAGE REJECT. In this case, both devices shall go to asynchronous data transfer mode for data transfers between the two devices. For the MESSAGE PARITY ERROR case, the implied agreement shall be reinstated if a retransmittal of the second of the pair of messages is successfully accomplished. After one retry attempt, if the target receives a MESSAGE PARITY ERROR message, it shall terminate the retry activity. This may be done either by changing to any other information transfer phase and transferring at least one byte of information or by going to the BUS FREE phase (see 3.1.1). The initiator shall accept such action as aborting the negotiation, and both devices shall go to asynchronous data transfer mode for data transfers between the two devices.

If the target recognizes that negotiation is required, it sends an SDTR message to the initiator. Prior to releasing the ACK signal on the last byte of the SDTR message from the target, the initiator shall assert the ATN signal and respond with its SDTR message or with a REJECT MESSAGE. If an abnormal condition prevents the initiator from returning an appropriate response, both devices shall go to asynchronous data transfer mode for data transfers between the two devices.

Following an initiator's responding SDTR message, an implied agreement for synchronous operation shall not be considered to exist until the target leaves the MESSAGE OUT PHASE, indicating that the target has accepted the negotiation. After one retry attempt, if the target has not received the initiator's responding SDTR message, it shall go to the BUS FREE phase without any further information transfer attempt (see 3.1.1). This indicates that a catastrophic error condition has occurred. Both devices shall go to asynchronous data transfer mode for data transfers between the two devices.

Note: SCSI devices capable of wide data transfers (>8 bit plus parity) shall initiate a WDTR message before initiating the SDTR message.

If, following an initiator's responding SDTR message, the target shifts to MESSAGE IN phase and the first message in is MESSAGE REJECT, the implied agreement shall be considered to be negated and both devices shall go to asynchronous data transfer mode for data transfers between the two devices.

The implied synchronous agreement shall remain in effect until a BUS DEVICE RESET message is received, until a hard reset condition occurs, or until one of the two SCSI devices elects to modify the agreement. The default data transfer mode is asynchronous data transfer mode. The default data transfer mode is entered at power on, after a BUS DEVICE RESET message, or after a hard reset condition.

Note: Renegotiation at every selection is not recommended, since a significant performance impact is likely.

The REQ/ACK offsets supported by a particular drive are given in a Table in that drive's individual Product Manual (Vol. 1).

Wide Data Transfer Request Message**Table 3.5.3-5. Wide data transfer message**

Bit Byte(s)	7	6	5	4	3	2	1	0
0	Extended message (01h)							
1	Extended message length (02h)							
2	Wide Data Transfer Request code (03h)							
3	Transfer Width Exponent							

A WIDE DATA TRANSFER REQUEST (WDTR) message (Table 3.5.3-5) exchange shall be initiated by an SCSI device whenever a previously-arranged transfer width agreement may have become invalid. The agreement becomes invalid after any condition which may leave the data transfer agreement in an indeterminate state such as:

- 1) after a hard reset condition;
- 2) after a BUS DEVICE RESET message and;
- 3) after a power cycle.

Note: For SCSI devices that implement WDTR, a WDTR negotiation invalidates a prior SDTR negotiation.

In addition, an SCSI device may initiate an WDTR message exchange whenever it is appropriate to negotiate a new transfer width agreement. SCSI devices that are capable of wide data transfers (greater than eight bits) shall not respond to an WDTR message with a MESSAGE REJECT message.

The WDTR message exchange establishes an agreement between two SCSI devices on the width of the data path to be 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 shall use an eight-bit data path.

If an SCSI device implements both wide data transfer option and synchronous data transfer option, then it shall negotiate the wide data transfer agreement prior to negotiating the synchronous data transfer agreement. If a synchronous data transfer agreement is in effect, then an SCSI device that accepts a WDTR message shall reset the synchronous agreement to asynchronous mode.

Byte 3 selects the transfer width in bytes. The number of bytes transfer width is 2^m bytes, where m is the Transfer Width Exponent given in byte 3. 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 ($m = 01h$). Values of m greater than 01h are not applicable to drives supported by this manual.

The originating SCSI device (the SCSI device that sends the first of the pair of WDTR messages) sets its transfer width value to the maximum data path width it elects to accommodate. 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 WDTR message. The successful completion of an exchange of WDTR messages implies an agreement as follows:

**Responding Device
WDTR Response****Implied Agreement**

- | | |
|----------------------------------|---|
| (1) Non-zero transfer width | Each device transmits and receives data with a transfer width equal to the responding SCSI device's transfer width. |
| (2) Transfer width equal to zero | Eight-bit Data transfer |
| (3) MESSAGE REJECT message | Eight-bit Data transfer |

If the initiator recognizes that negotiation is required, it asserts the ATN signal and sends a WDTR message to begin the negotiating process. After successfully completing the MESSAGE OUT phase, the target shall respond with the proper WDTR message. If an abnormal condition prevents the target from returning an appropriate response, both devices shall go to eight-bit data transfer mode for data transfers between the two devices.

Following target response (1) above, the implied agreement for wide data transfers shall be considered to be negated by both the initiator and the target if the initiator asserts ATN and the first message out is either MESSAGE PARITY ERROR or MESSAGE REJECT. In this case, both devices shall go to eight-bit data transfer mode for data transfers between the two devices. For the MESSAGE PARITY ERROR case, the implied agreement shall be reinstated if a re-transmittal of the second of the pair of messages is successfully accomplished. After a vendor-specific number of retry attempts (greater than zero), if the target receives a MESSAGE PARITY ERROR message, it shall terminate the retry activity. This may be done either by changing to any other information transfer phase and transferring at least one byte of information or by going to the BUS FREE phase (see 3.1.1). The initiator shall accept such action as aborting the negotiation, and both devices shall go to eight-bit data transfer mode for data transfers between the two devices.

If the target recognizes that negotiation is required, it sends a WDTR message to the initiator. Prior to releasing the ACK signal on the last byte of the WDTR message from the target, the initiator shall assert the ATN signal and respond with its WDTR message or with a MESSAGE REJECT message. If an abnormal condition prevents the initiator from returning an appropriate response, both devices shall go to eight-bit data transfer mode for data transfers between the two devices.

Following an initiator's responding WDTR message, an implied agreement for wide data transfer operation shall not be considered to exist until the target leaves the MESSAGE OUT phase, indicating that the target has accepted the negotiation. After a vendor-specific number of retry attempts (greater than zero), if the target has not received the initiator's responding WDTR message, it shall go to the BUS FREE phase without any further information transfer attempt (see 3.1.1). This indicates that a catastrophic error condition has occurred. Both devices shall go to eight-bit data transfer mode for data transfers between the two devices.

If, following an initiator's responding WDTR message, the target shifts to MESSAGE IN phase and the first message in is MESSAGE REJECT, the implied agreement shall be considered to be negated and both devices shall go to eight-bit data transfer mode for data transfers between the two devices.

The implied transfer width agreement shall remain in effect until a BUS DEVICE RESET message is received, until a hard reset condition occurs, or until one of the two SCSI devices elects to modify the agreement. The default data transfer width is eight-bit data transfer mode. The default data transfer mode is entered at power on, after a BUS DEVICE RESET message, or after a hard reset condition.

16-bit Wide data Transfer

Wide data transfer is optional and may be used in the DATA phase only if a nonzero wide data transfer agreement is in effect. The messages determine the use of wide mode by both SCSI devices and establish a data path width to be used during the DATA phase. Default for the drives that have wide data transfer capability is 8-bit data transfers.

During 16-bit wide data transfers, the first and second logical data bytes for each data phase shall be transferred across the DB(7-0,P) and DB(15-8,P1) signals respectively, on the 68 conductor cable used by drives that support the wide data transfer feature. Subsequent pairs of data bytes are likewise transferred in parallel across the 68 conductor cable.

To illustrate the order of transferring data bytes across the interface assume an example four byte data transfer of bytes W, X, Y and Z.

When transferring bytes W, X, Y and Z across the 8 or 16-bit bus widths, they are transferred in the order shown below:

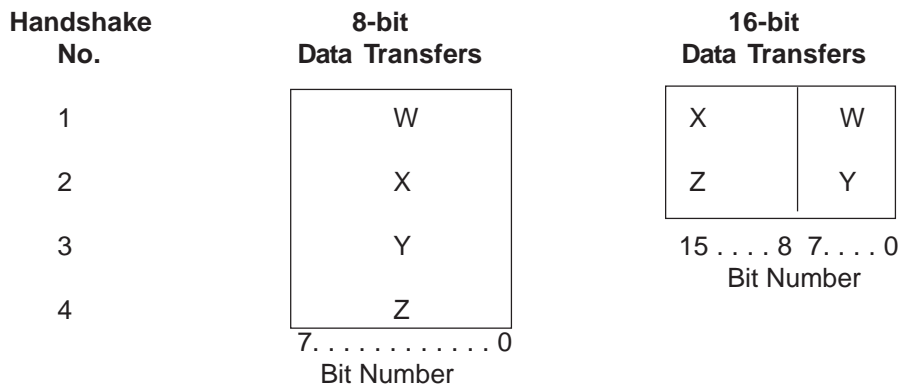


Figure 3.5.3-1. Wide SCSI byte ordering

If the last data byte transferred for a command does not fall on the DB(15-8,P1) signals for a 16-bit wide transfer, then the values of the remaining higher-numbered bits are undefined. However, parity bits for these undefined bytes shall be valid for whatever data is placed on the bus.

Ignore Wide Residue (23h)

The Ignore Wide Residue message is not an extended message, but it is placed here in context with the 16-bit Wide Data Transfer message, since it applies when the 16-bit wide data transfer capability is used.

Table 3.5.3-6. Ignore wide residue message format

Bit Byte(s)	7	6	5	4	3	2	1	0
0	Message Code (23h)							
1	Ignore (01h, 02h, 03h)							

The IGNORE WIDE RESIDUE message (Table 3.5.3-6) shall be sent from a target to indicate that the number of valid bytes sent during the last REQ/ACK handshake of a DATA IN phase is less than the negotiated transfer width. The ignore field indicates the number of invalid data bytes transferred. This message shall be sent immediately following that DATA IN phase and prior to any other messages. The ignore field is defined in Table 3.5.3-7.

Note. More than one IGNORE WIDE RESIDUE message may occur during an I/O process.

Table 3.5.3-7. Ignore field definition

Ignore	Invalid Data Bits 16-bit Transfers
00h	Reserved
01h	DB(15-8)
02h	Reserved
03h	Reserved
04h - FFh	Reserved

Even though a byte is invalid its corresponding parity bit shall be valid for the value transferred.

3.6 Message exception conditions handling

Previous versions of this SCSI I/O Product Manual were not clear on how message exception conditions should be handled. This section with its accompanying Table 3.6.1-1 attempts to clarify how to handle message exception conditions. The numbers in the grid spaces of Table 3.6.1-1 are response code numbers that are decoded and explained in the notes following the table. The abbreviations along the top of the table are expanded at the end of the table also.

Table 3.6.1-1. SCSI Message Handling

MESSAGE	PHASE	Sel	ID	Mout	Cmd	M-in	Resel	Disc	Data	Stat	Cplt
ABORT (06)		2	2	2	2	2	2	2	2	2	2
ABORT TAG (0D)		4	2	2	2	2	2	2	2	2	2
BUS DEVICE RESET (0C)		2	2	2	2	2	2	2	2	2	2
BUS DEV RST OTHER PORTS (14)*		4	1	1	1	1	1	9,1	1	1	9,1
CLEAR QUEUE (0E)		4	2	2	2	2	2	2	2	2	2
CONTINUE I/O PROCESS (12)		4	1	1	3,1	3,1	3,1	3,9,1	3,1	3,1	3,9,1
INITIATOR DETECTED ERR (05)		4	5	7	6	5	5	5	7	6	5
IDENTIFY (Invalid - changed the value of LUNTRN.)		N/A	C	C	C	C	C	C	C	C	C
IDENTIFY (Invalid - LUNTR bit or a reserved bit is set)		3,4	3,1	3,1	3,1	3,1	3,1	3,9,1	3,1	3,1	3,9,1
IDENTIFY (Valid)**		1	1	1	1	1	1	9,1	1	1	9,1
MESSAGE PARITY ERROR (09)		4	4	C	C	5	5	5	C	C	5
MESSAGE REJECT (07)		4	3,1	3,1	3,1	A	C	8	3,1	3,1	9,1
NO OP (08)		4	1	1	1	1	1	9,1	1	1	9,1
QUEUE - SIMPLE (20)		4	1	3,1	3,1	3,1	3,1	3,9,1	3,1	3,1	3,9,1
QUEUE - ORDERED (22)		4	1	3,1	3,1	3,1	3,1	3,9,1	3,1	3,1	3,9,1
QUEUE - HEAD (21)		1,4	1	3,1	3,1	3,1	3,1	3,9,1	3,1	3,1	3,9,1
SYNCHRONOUS TRANSFER REQ		4	1	1	1	1	1	9,1	1	1	9,1
TARGET TRANS DIS (13)		4	1	1	3,1	3,1	3,1	3,9,1	3,1	3,1	3,9,1
WIDE TRANSFER REQUEST		4	1	1	1	1	1	9,1	1	1	9,1
Invalid, unimplemented, illegal msgs		4	3,1	3,1	3,1	3,1	3,1	3,9,1	3,1	3,1	3,9,1
Parity error detected by target		B	5	5	6	N/A	N/A	N/A	7	N/A	N/A

* The responses for the Bus Device Reset Other Ports message assume that the drive will implement this message. When the message isn't implemented (i.e. on drives which do not support dual-port SCSI), the responses will be as stated in the "Invalid, unimplemented, illegal msgs" row.

** In a multi-initiator or queued environment, an initiator's attempt to prevent disconnection (by clearing the DiscPriv bit in the initial Identify message or in a subsequent Identify message, or by rejecting a Disconnect message) may result in the target terminating the associated command with a status of Busy.

PHASE/MESSAGE**ABBREVIATIONS**

SELECTION	Sel
IDENTIFY (AFTER SELECTION)	ID
MSG OUT	Mout
COMMAND	Cmd
MSG IN (NOT CPLT OR DISC).	M-in
RESELECT (AFTER IDENTIFY)	Resel
MSG IN (DISCONNECT)	Disc
DATA IN/OUT	Data
STATUS	Stat
MSG IN (COMMAND COMPLETE)	Cplt

RESPONSES BY NUMBER

- 1: Continue. (See **note on previous page for multi-initiator or queued environment considerations)
- 2: Enter Bus Free state. No sense data will be set up.
- 3: Send Message Reject message.
- 4: "Unexpected Bus Free" without sense data set up.
- 5: Retry Message phase.
- 6: Send Restore Pointers message, and retry phase.
- 7: Go to Status phase and report Check Condition status. (See ** note on previous page for multi-initiator or queued environment considerations).
- 8: Continue with no disconnect. (See ** note on previous page for multi-initiator or queued environment considerations)
- 9: Re-send message (Disconnect or Command Complete).
- A: If Message In is Save Data Pointer, do response 8. Otherwise, do response 1.
- B: Do not respond to the selection.
- C: "Unexpected Bus Free" with sense data set up.
- N/A: Not applicable

NOTES RELATIVE TO RESPONSES

(Response #) Additional Information

- (1) The Continue response indicates that the target will continue execution of the current command. If the message sent to the target is a wide or a synchronous negotiation request, the target will complete the negotiation prior to continuing the command.
- (4) The "Unexpected Bus Free" response includes going directly to bus free (ATN signal is ignored if present) and terminating the current command. Since no sense data will be set up, a Contingent Allegiance condition will not be established.
- (5) On a Message Out phase, the target will continue to receive message bytes from the initiator until the ATN signal is deserted. The target will then remain in the Message Out phase and assert REQ, to notify the initiator that the entire Message Out phase needs to be repeated. Following failure of a product-specific number of retries, the target will switch to response C.

On a Message In phase, the target will resend the last message. Following failure of a product-specific number of retries, the target will switch to response C.
- (7) Associated sense data will have a sense key of ABORTED COMMAND (0Bh). The sense code will either indicate an Initiator Detected Error message was received or that a Parity Error was detected, as applicable.
- (8) This is a special case of (1). The target should continue the I/O process without releasing the bus. The target may attempt to disconnect at a later time.
- (9) If the message sent to the target is a wide or a synchronous negotiation request, the target will complete the negotiation prior to resending the command complete or the disconnect message.
- (C) The "Unexpected Bus Free" response includes going directly to Bus Free (ATN signal is ignored if present) and terminating the current command. The sense data will be set up for the failing initiator with a sense key of ABORTED COMMAND (0Bh) and a sense code of INVALID MESSAGE ERROR (49h). A Contingent Allegiance condition will be established.
Exceptions:
 - 1) If the "Unexpected Bus Free" occurs due to a message retry failure (response 5), the sense code will be set to either indicate an Initiator Detected Error message was received or that a Parity Error was detected, as applicable.

3.7 S.M.A.R.T. system

Some drive families mentioned in Section 1.1 implement what is called in the industry the S.M.A.R.T. system. S.M.A.R.T. is an acronym for Self-Monitoring Analysis and Reporting Technology. The intent of the S.M.A.R.T. system is to recognize conditions that indicate imminent drive failure, and provide sufficient warning to the host system of impending failure. The host system may use the information provided to trigger it to perform diagnostic, preventative, and/or protective functions (e.g. data backup).

The initiator sets up the parameters for S.M.A.R.T. operation using Mode Select Informational Exceptions Control page 1Ch. The drive reports information about S.M.A.R.T. operation using Request Sense Additional Sense Code 5D 00 and Mode Sense data page 1Ch. Refer to Section 5.2.1 for the description of the Mode Select/Mode Sense commands and more details on the Informational Exceptions Control page.

Refer to the Volume 1 Product Manual for each drive family to determine the extent of its implementation of the S.M.A.R.T. system.

4.0 SCSI commands

This section defines the SCSI command structure and describes a typical SCSI bus procedure involving a command, status return and message interchange.

The command structure defined herein provides for a contiguous set of logical blocks of data to be transferred across the interface. The number of logical data blocks to be transferred is defined in the command. Initiator commands to the drive are structured in accordance with the requirements imposed by the drive physical characteristics. These physical characteristics are reported to the initiator in response to an inquiry command.

A single command may transfer one or more logical blocks of data. The drive may disconnect, if allowed by the initiators, from the SCSI bus to allow activity by other SCSI devices while the drive performs operations within itself.

Upon command completion (which may be executed either successfully or unsuccessfully), the drive returns a status byte to the initiator. Since most error and exception conditions cannot be adequately described with a single status byte, one status code that can be sent as the status byte is called Check Condition. It indicates that additional information is available. The initiator may issue a Request Sense command to request the return of the additional information as part of the Data In phase of the command.

4.1 Command implementation requirements

The first byte of any SCSI command contains an operation code as defined in this document. Three bits (bits 7-5) of the second byte of each SCSI command specify the logical unit if it is not specified using the Identify Message (see Paragraph 3.5.3.1). Only logic unit zero is valid for drives to which this specification applies. The last byte of all SCSI commands shall contain a control byte as defined in Paragraph 4.2.6.

4.1.1 Reserved

Reserved bits, bytes, fields, and code values are set aside for future standardization. Their use and interpretation may be specified by future revisions to this specification. A reserved bit, field, or byte shall be set to zero, or in accordance with a future revision to this specification. A drive that receives a reserved code value shall terminate the command with a Check Condition status and the Sense Key shall be set to Illegal Request. It shall also be acceptable for the drive to interpret the bit, field, byte, or code value in accordance with a future revision to this specification.

4.2 Command descriptor block (CDB)

A request by an initiator to a drive is performed by sending a Command Descriptor Block (CDB) to the drive. For several commands, the request is accompanied by a list of parameters sent during the Data Out phase. See the specific commands for detailed information.

The Command Descriptor Block always has an operation code as the first byte of the command. This is followed by a logical unit number, command parameters (if any), and a control byte.

For all commands, if there is an invalid parameter in the Command Descriptor Block, the drive shall terminate the command without altering the medium.

The format description for the Command Descriptor Block as supported by the drive is shown in Tables 4.2-1, 4.2-2 and 4.2-3.

4.2.1 Operation code

The operation code (Table 4.2-1) of the Command descriptor Block has a group code field and a command code field. The three bit group code field provides for eight groups of command codes. The five bit command code field provides for thirty two command codes in each group. Thus, a total of 256 possible operation codes exist. Operation codes are defined in Section 5.

For the drive the group code specifies one of the following groups:

Group 0 - Six byte commands (see Table 4.2-2).

Group 1 - Ten byte commands (see Table 4.2-3).

Group 2 - Ten byte commands (see Table 4.2-3).

Group 3 - Reserved

Group 4 - Reserved

Group 5 - Twelve-byte commands

Group 6 - Vendor specific

Group 7 - Vendor Specific

Table 4.2-1. Operation Code Format for CDB

Bit Byte(s)	7	6	5	4	3	2	1	0
0	Group Code			Command Code				

Table 4.2-2. Typical Command Descriptor Block for Six Byte Commands

Bit Byte(s)	7	6	5	4	3	2	1	0
0	Operation Code							
1	0	0	0	Logical Block Address (if req.) (MSB)				
2	Logical Block Address (if required)							
3	Logical Block Address (if required)(LSB)							
4	Transfer Length (if required)							
5	Control Byte							

Table 4.2-3. Typical Command Descriptor Block for Ten Byte Commands

Bit Byte(s)	7	6	5	4	3	2	1	0
0	Operation Code							
1	Logical Unit No. 0 0 0			Reserved				Rel Adr
2	Logical Block Address (if required) (MSB)							
3	Logical Block Address (if required)							
4	Logical Block Address (if required)							
5	Logical Block Address (if required) (LSB)							
6	Reserved							
7	Transfer Length (if required) (MSB) [1]							
8	Transfer Length (if required) (LSB) [1]							
9	Control Byte [1]							

[1] For Twelve byte commands there are two more bytes of transfer length, bytes 9 and 10, and the control byte is byte 11.

4.2.2 Logical Unit Number (LUN)

The logical unit number (LUN) addresses one of up to eight physical devices or virtual devices attached to a target. The only valid LUN number for the drives supported by this manual is Zero.

The LUN in the CDB is provided for systems that do not implement the Identify Message. If an Identify message is sent to the drive, the drive will use the LUN number specified in this message. In this case, the drive shall ignore the LUN specified within the command descriptor block.

The drive will reject commands that select an invalid LUN (except Request Sense and Inquiry) by requesting and accepting the command bytes, then going to Status phase and sending Check Condition status. Note that the LUN is sent in the LUN field of a CDB (if no Identify message has been received for this selection) or by the LUN field of an Identify message.

Request Sense commands selecting an invalid LUN will receive a Sense Data block with the Illegal Request Sense Key and an Invalid LUN Error Code. Inquiry commands will return Inquiry Data with the Peripheral Device Type field set to Logical Unit Not Present (7Fh). Request Sense and Inquiry commands will not send Check Condition status in response to an invalid LUN selection.

4.2.3 Logical Block Address

The logical block address in the command descriptor block shall begin with block zero and be continuous up to the last logical block on drive.

Group 0 command descriptor block contains 21 bit logical block addresses. Groups 1 & 2 command descriptor blocks contain 32 bit logical block addresses.

The logical block concept implies that the initiator and target shall have previously established the number of data bytes per logical block. This may be established through the use of the Read Capacity command or the Mode Sense command or by prior arrangement.

The maximum logical block address for the drive which is accessible by the Initiator is defined in Read Capacity Command data in section 5.2.2.1.

4.2.4 Relative address bit

Relative addressing is a technique useful in accessing structured data in a uniform manner. Relative addressing is only allowed when commands are linked. Details are given with those commands that use this feature.

4.2.5 Transfer length

The Transfer Length specifies the amount of data to be transferred, usually the number of blocks. For several commands the transfer length indicates the requested number of bytes to be sent as defined in the command description. For these commands the transfer length field may be identified by a different name. See the following descriptions and the individual command descriptions for further information.

Commands that use one byte for Transfer Length allow up to 256 blocks of data to be transferred by one command. A Transfer Length value of 1 to 255 indicates the number of blocks that shall be transferred. A value of zero indicates 256 blocks.

Commands that use two bytes for Transfer Length allow up to 65,535 blocks of data to be transferred by one command. In this case, a Transfer Length of zero indicates that no data transfer shall take place. A value of 1 to 65,535 indicates the number of blocks that shall be transferred.

For several commands more than two bytes are allocated for Transfer Length. Refer to the specific command description for further information.

The Transfer Length of the commands that are used to send a list of parameters to a drive is called the Parameter List Length. The Parameter List Length specifies the number of bytes sent during the Data Out phase.

The Transfer Length of the commands used to return sense data (e.g. Request Sense, Inquiry, Mode Sense, etc.) to an initiator is called the Allocation Length. The Allocation Length specifies the number of bytes that the initiator has allocated for returned data. The drive shall terminate the Data In phase when Allocation Length bytes have been transferred or when all available data have been transferred to the initiator, whichever is less.

4.2.6 Control byte

Normally all zeros. Typically, the drive does not support the Flag and Link bit functions. It can be supported as a special factory installed option.

Table 4.2.6-1. Control byte

Bit Byte(s)	7	6	5	4	3	2	1	0
LAST	0	0	0	0	0	0	Flag* 0 or 1	Link* 0 or 1

*Bit	Description
1	Flag bit - If the link bit is zero, then the flag bit shall be set to zero. If the link bit is one, and if the command terminates successfully, the target shall send LINKED COMMAND COMPLETE message if the flag bit is zero and shall send LINKED COMMAND COMPLETE (WITH FLAG) message if the flag bit is one. Typically, this bit is used to cause an interrupt in the initiator between linked commands.
0	Link bit - This bit is set to one to indicate that the initiator desires an automatic link to the next command upon successful completion of the current command. Implementation of linked commands is optional. If the link bit is one, upon successful termination of the command, the target shall return INTERMEDIATE status and shall then send one of the two messages defined by the flag bit (above).

4.3 Status

A status byte shall be sent from the target to the initiator during the STATUS phase at the termination of each command as specified in Tables 4.3-1 and 4.3-2 unless the command is cleared by one of the following conditions:

1. an ABORT message
2. a BUS DEVICE RESET message
3. a hard reset condition
4. an unexpected BUS FREE condition (see 3.1.1)
5. an ABORT TAG message
6. a CLEAR QUEUE message

Table 4.3-1. Status byte

Bit Byte(s)	7	6	5	4	3	2	1	0
0	Reserved		Status Byte Code					Rsvd

Table 4.3-2. Status byte code bit values

<u>Bits of Status Byte</u>								
7	6	5	4	3	2	1	0	Status Represented
R	R	0	0	0	0	0	R	Good
R	R	0	0	0	0	1	R	Check Condition
R	R	0	0	0	1	0	R	Condition Met/Good
R	R	0	0	1	0	0	R	Busy
R	R	0	1	0	0	0	R	Intermediate/Good
R	R	0	1	0	1	0	R	Intermediate/Condition Met
R	R	0	1	1	0	0	R	Reservation Conflict
R	R	1	0	0	0	1	R	Command Terminated
R	R	1	0	1	0	0	R	Queue Full
All Other Codes								Reserved

Key: R - Reserved bit (must be zero for drive)

A description of the status byte codes is given below.

Good - This status indicates that the target has successfully completed the command.

Check Condition - Any error, exception, or abnormal condition that causes sense data to be set, shall cause a CHECK CONDITION status. The REQUEST SENSE command should be issued following a CHECK CONDITION status, to determine the nature of the condition.

Condition Met - The SEARCH DATA commands shall return this status whenever a search condition is satisfied. This status does not break a chain of linked commands. The logical block address of the logical block that satisfies the search may be determined with a REQUEST SENSE command.

Busy - The target is busy. This status shall be returned whenever a target is unable to process the command from an otherwise acceptable initiator. BUSY status shall be returned if the initiator has not granted the disconnect privilege and attempts to queue a command or if the initiator rejects the disconnect message and the queue is not empty. BUSY status shall also be returned if a Change Definition command is received requesting that the drive change from SCSI-2 mode to SCSI-1 mode and the queue is not empty, or if commands from other initiators or tagged commands from the contingent host are received while a Contingent Allegiance condition exists. The normal initiator recovery action is to issue the command again at a later time.

Intermediate - This status will be returned for every command in a series of linked commands (except the last command), unless an error, exception or abnormal condition causes a CHECK CONDITION status, a RESERVATION CONFLICT STATUS or a COMMAND TERMINATED status to be set. If this status is not returned, the chain of linked commands is broken; no further commands in the series are executed.

Reservation Conflict - This status shall be returned whenever a SCSI device attempts to access a logical unit or an extent within a logical unit that is reserved with a conflicting reservation type for another SCSI device (see RESERVE and RESERVE UNIT command). The normal initiator recovery action is to issue the command again at a later time.

Command Terminated - This status shall be returned whenever the target terminates the current I/O process after receiving a TERMINATE I/O PROCESS message (see 3.5.2). This status also indicates that a contingent allegiance condition has occurred (see 3.2.3). This message is not supported on standard OEM drives, but is a factory installed option.

Queue Full - This status shall be implemented if tagged queuing is implemented. This status is returned when a command is received and the command can not be accepted because the command queue is full. The command is not executed.

4.4 Command examples

4.4.1 Single command example

A typical operation on the SCSI bus is likely to include a single Read command to a peripheral device such as the drive. This operation is described in detail starting with a request from the initiator. This example assumes that no linked commands and no malfunctions or errors occur and is illustrated in Figure 4.4-1.

The initiator has active pointers and a set of stored pointers representing active disconnected SCSI devices (an initiator without disconnect capability does not require stored pointers). The initiator sets up the active pointers for the operation requested, arbitrates for the SCSI bus, and selects the drive. Once this process is completed, the drive assumes control of the operation.

The drive obtains the command from the initiator (in this case, a Read command). The drive interprets the command and executes it. For this command, the drive reads the requested data from the Disc Media and sends this data to the initiator. After sending the read data to the initiator, the drive sends a status byte to the initiator. To end the operation, the drive sends a Command Complete message to the initiator and then goes to the Bus Free state.

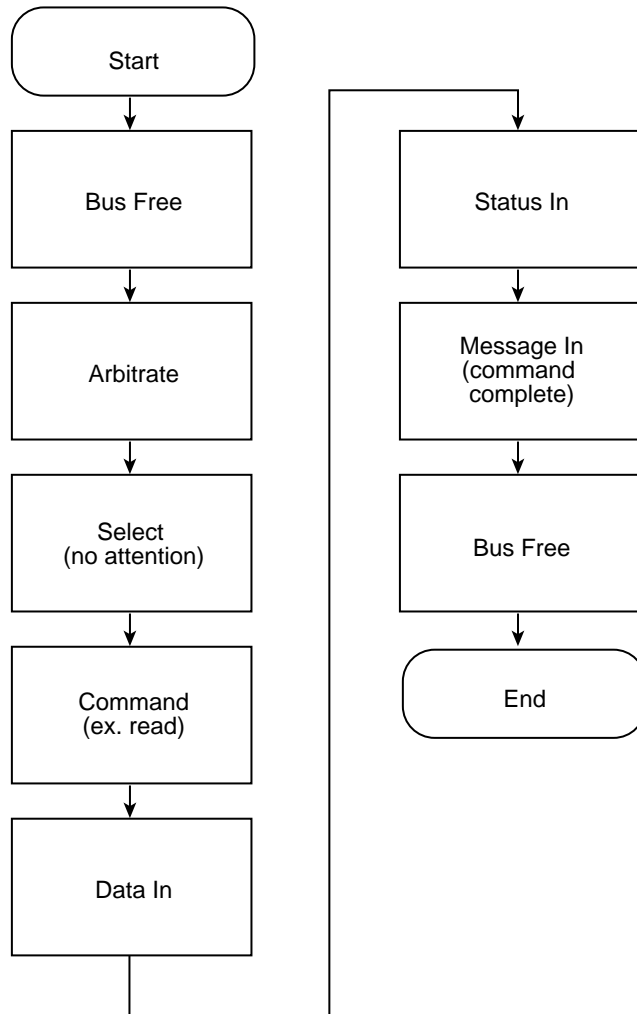


Figure 4.4-1. Single command example

4.4.2 Disconnect example

In the single command example, the length of time necessary to obtain the data may require a time consuming physical seek. In order to improve system throughput, the drive may disconnect from the initiator, freeing the SCSI bus to allow other requests to be sent to other SCSI devices. To do this, the initiator must be reselectable and capable of restoring the pointers upon reconnection. The drive must be capable of arbitrating for the SCSI bus and reselecting the initiator. See Figure 4.4-2.

After the drive has received the Read command (and has determined that there will be a delay), it disconnects by sending a Disconnect message and releasing BSY (goes to BUS Free state).

When the data is ready to be transferred the drive reconnects to the initiator, the initiator restores the pointers to their most recently saved values (which, in this case, are the initial values) and the drive continues (as in the single command example) to finish the operation. The initiator recognizes that the operation is complete when a Command Complete message is received.

If the drive elects to disconnect after transferring part of the data (e.g. while crossing a cylinder boundary), it sends a Save Data Pointer message and a Disconnect message to the initiator and then disconnects. When reconnection is completed, the initiator restores the current data pointer to the value it was immediately before the Save Data Pointer message.





Figure 4.4-2. Disconnect example

4.5 Timing examples

Times (T00 through T35) necessary to define performance are listed in the product manuals for each individual drive. Timing waveforms to define these times are illustrated in Tables 4.5-1 through 4.5-15.

Table 4.5-1. Arbitration, Selection, (No ATN), and Command Phase phase



* and 8-15 if applicable

Table 4.5-2. Arbitration, Selection (with ATN), and Message Out


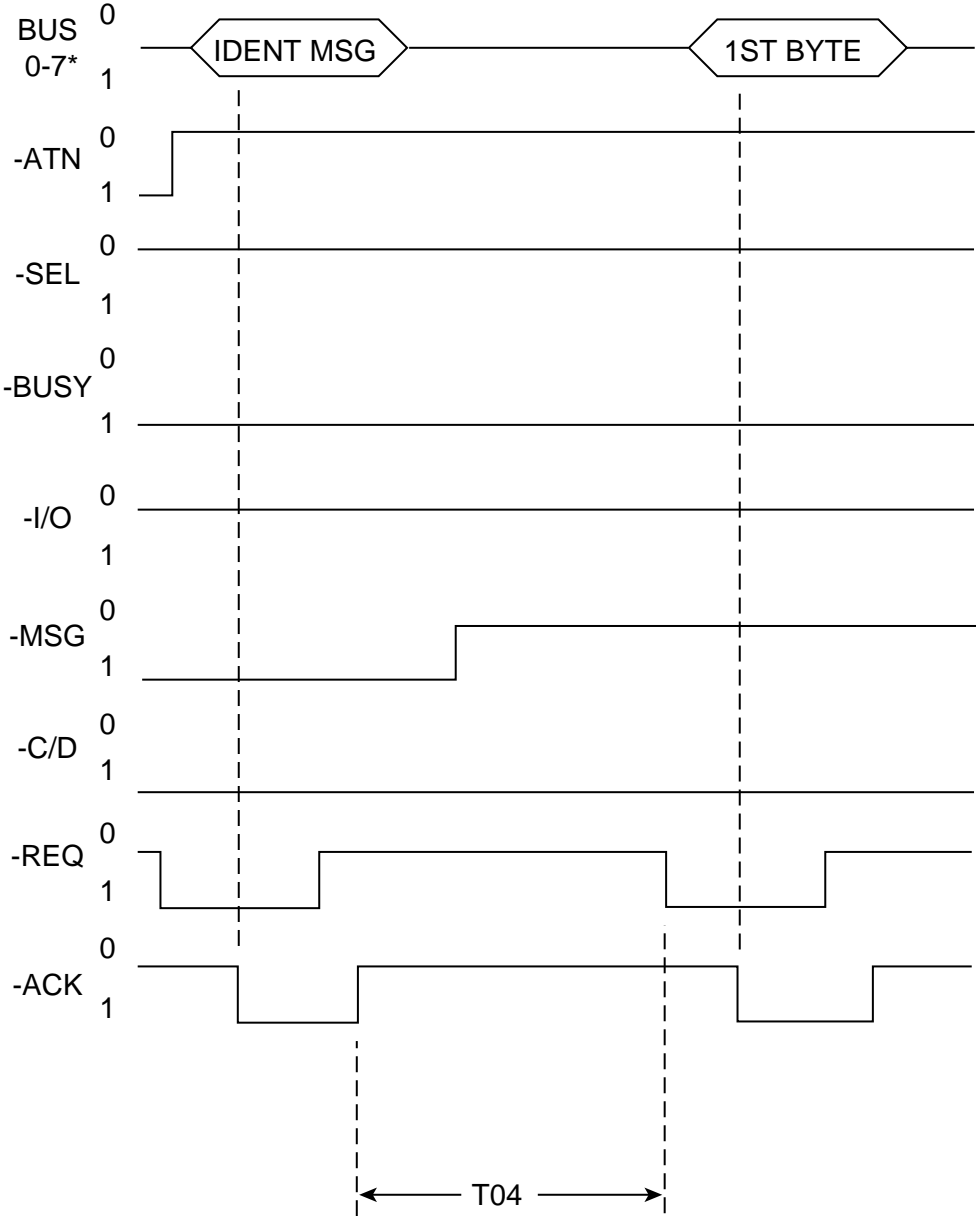
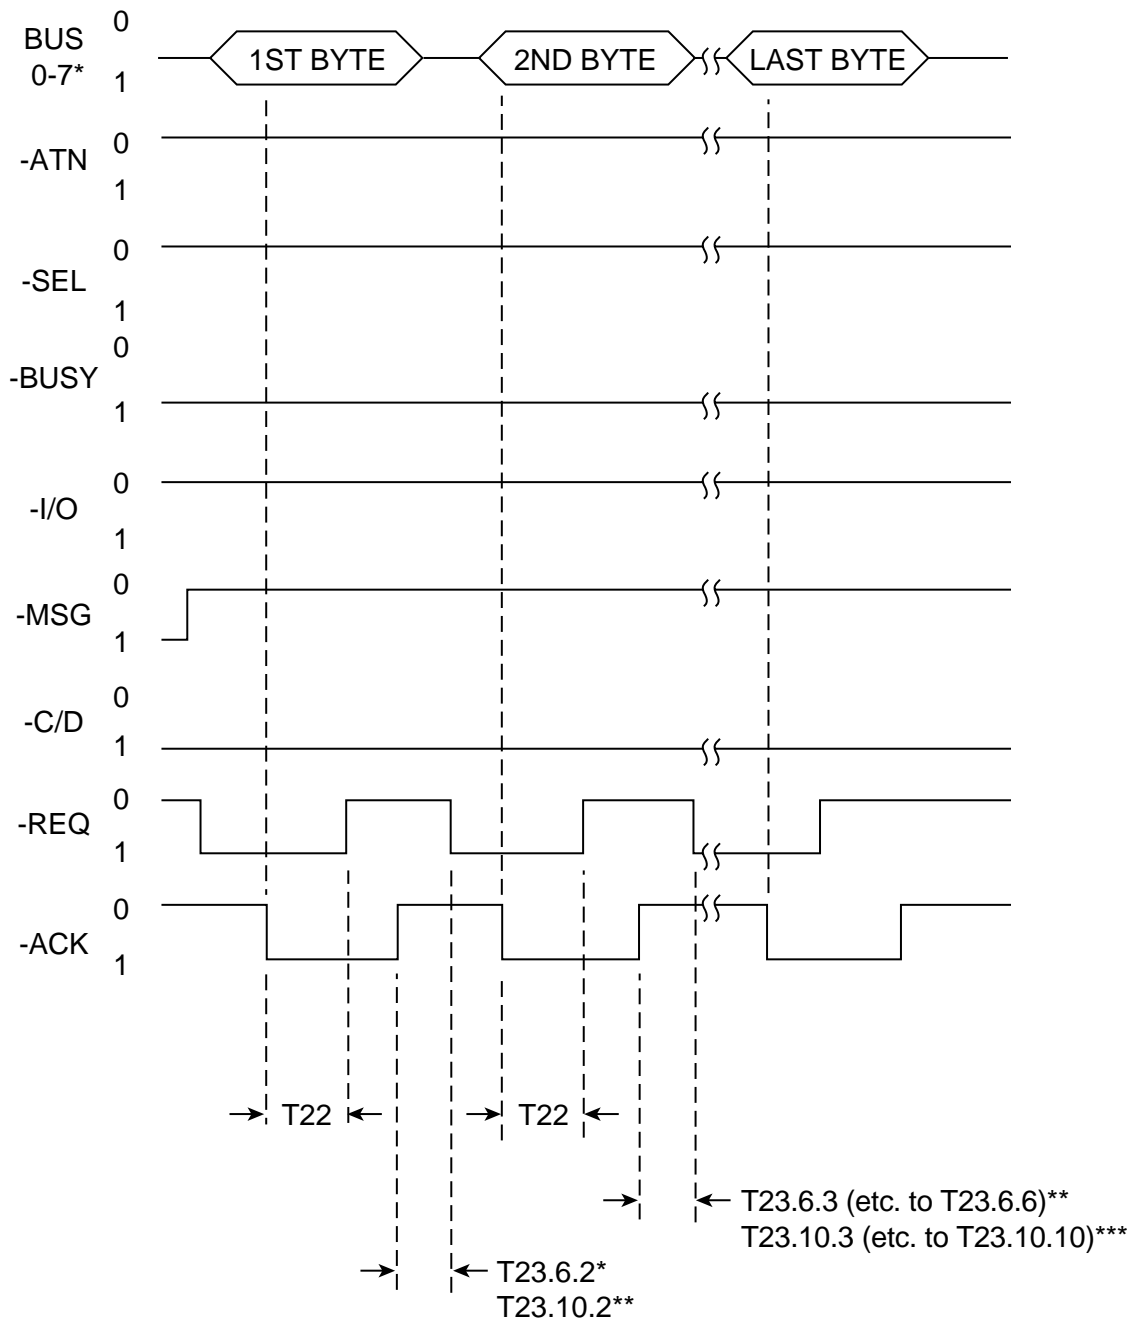


Table 4.5-3. Identify Msg Out to Command Phase



* and 8-15 if applicable

Table 4.5-4. Command Descriptor Block Transfer

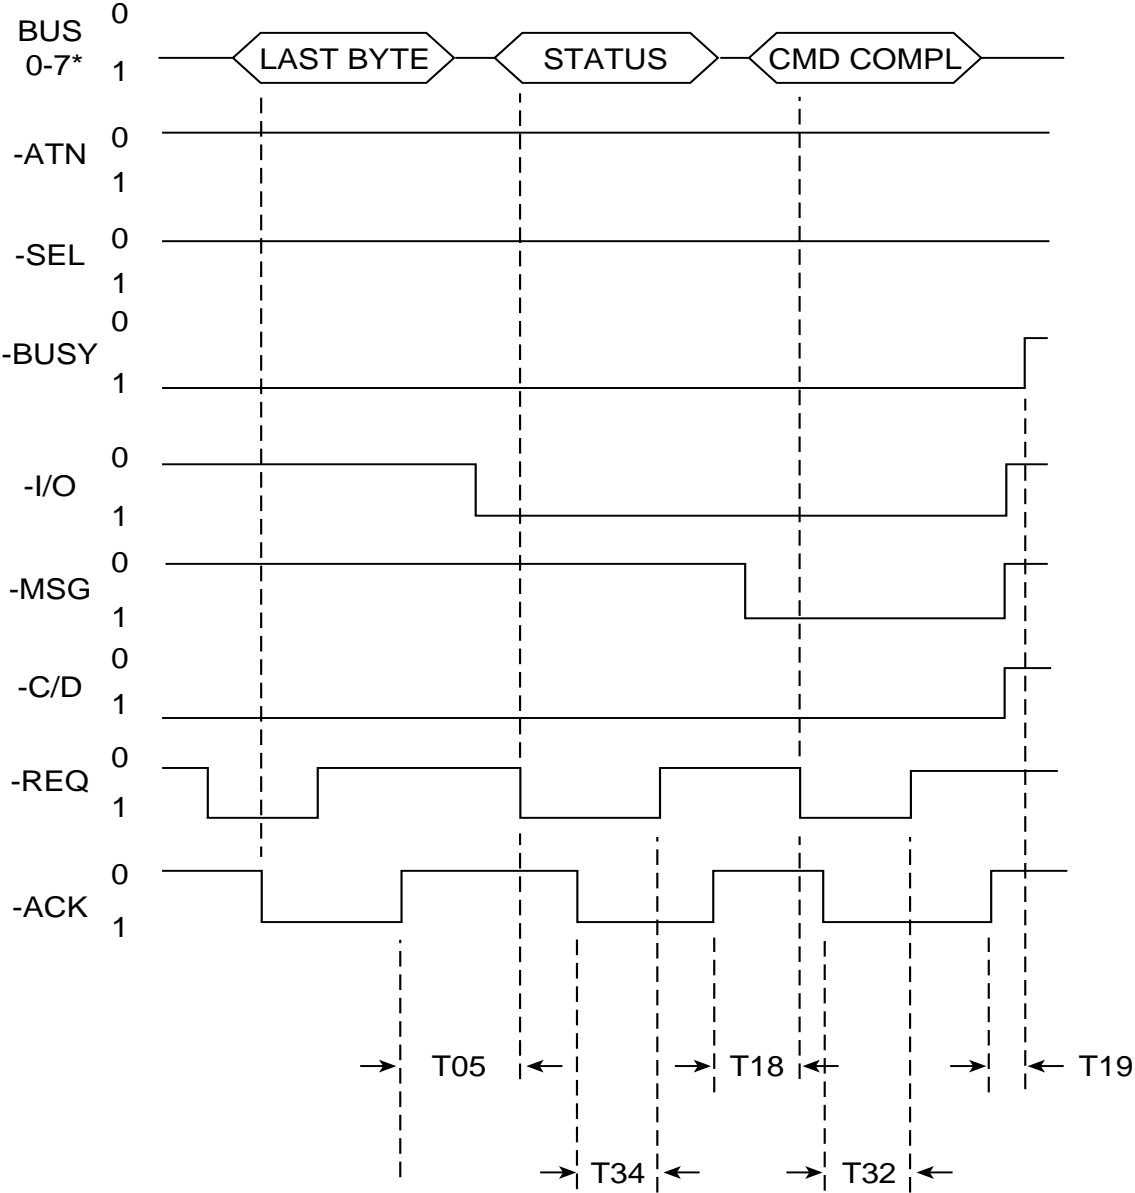


* and 8-15 if applicable

** six byte commands

*** ten byte commands

Table 4.5-5. Command Phase, Status Phase, Command Completed Msg and Bus Free



* and 8-15 if applicable

Table 4.5-6. Last Command Byte, Disconnect Msg, Bus Free, and Reselect





Table 4.5-7. Arbitration, Reselection and Message In



* and 8-15 if applicable

Table 4.5-8. Reselect Identify Msg, Status Phase, Command Complete Msg and Bus Free



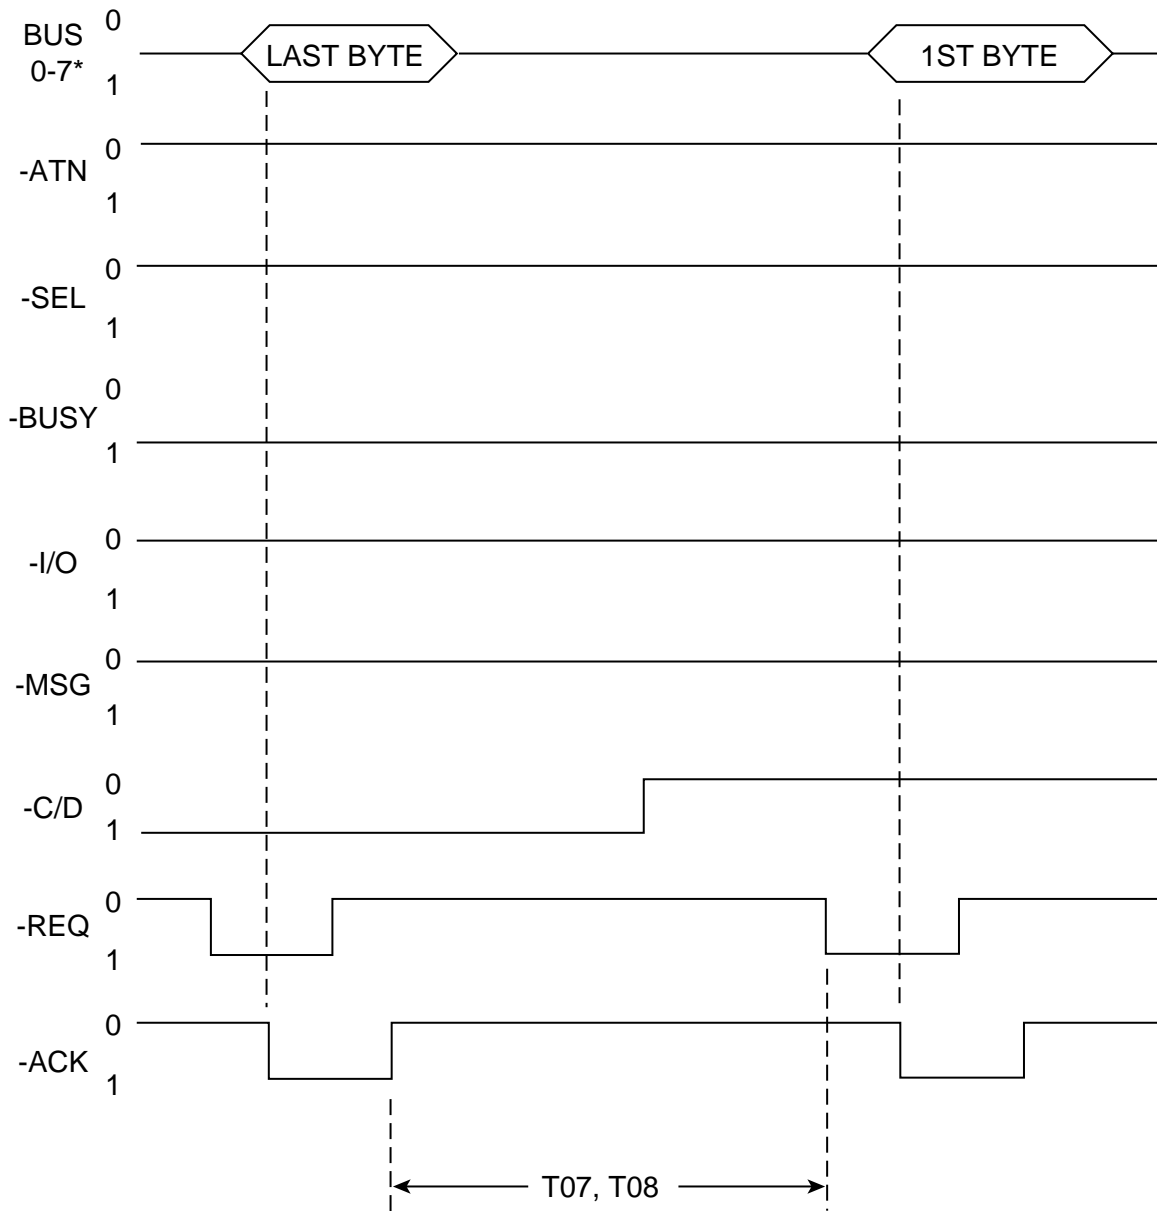


Table 4.5-9. Last Command Byte to Data in Phase




* and 8-15 if applicable

Table 4.5-10. Last Command Byte to Data Out Phase



* and 8-15 if applicable

Table 4.5-11. Reselect Identify Msg to Data in Phase



* and 8-15 if applicable

Table 4.5-12. Data in Block Transfer

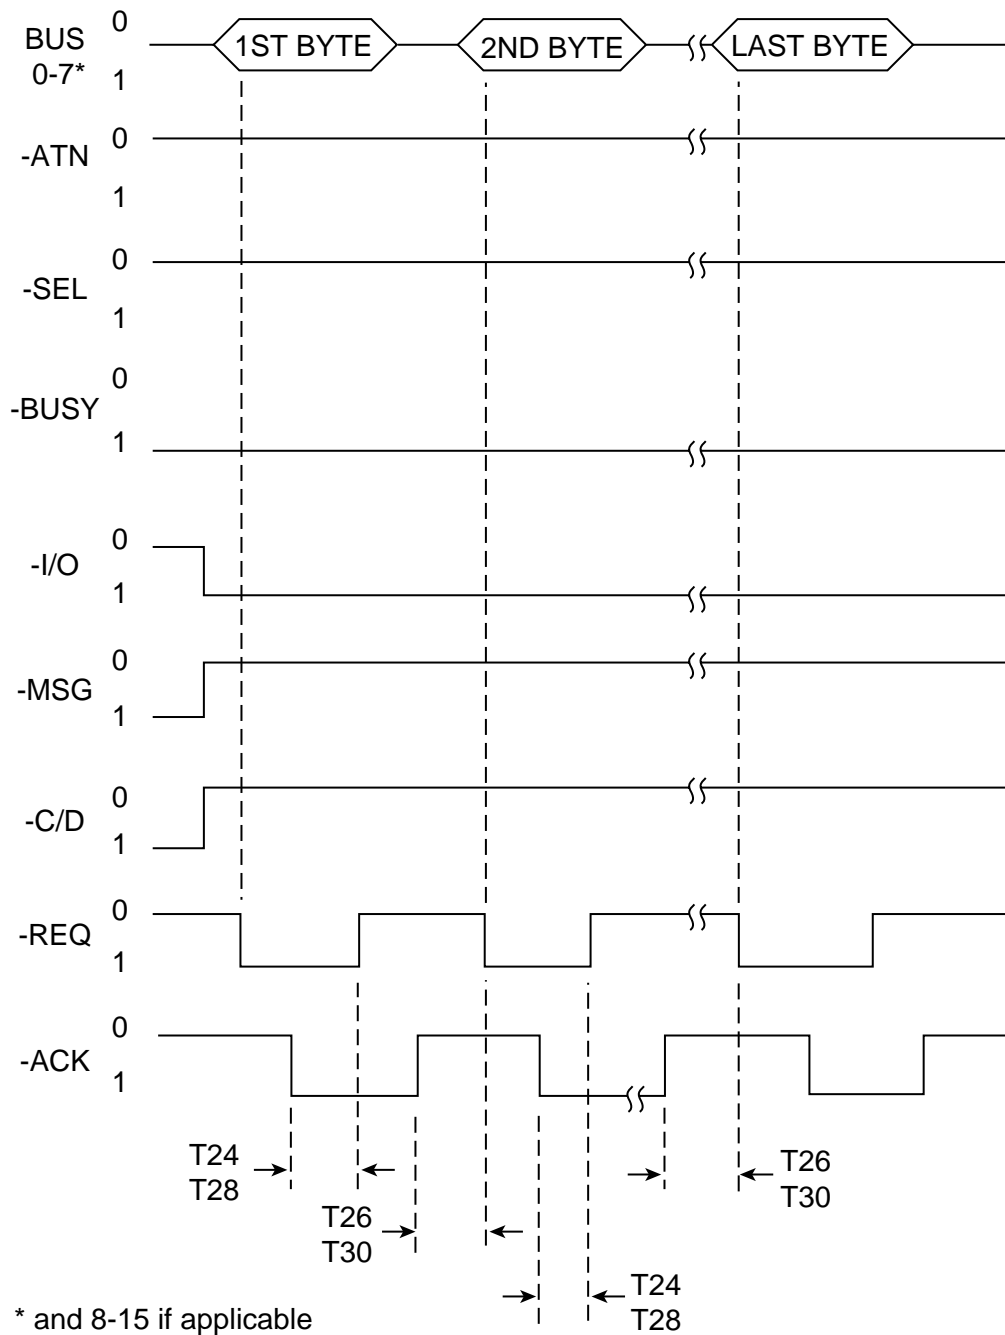




Table 4.5-13. Data Out Block Transfer




* and 8-15 if applicable

Table 4.5-14. Last Data Byte, Save Pointer Msg, and Disconnect Msg



* and 8-15 if applicible

Table 4.5-15. Data in Phase, Status Phase, Command Complete Msg, and Bus Free



* and 8-15 if applicable

4.6 Unit attention condition

The drive sets up the Unit Attention condition when it stores (within itself) a Unit Attention condition flag for each device on the SCSI bus having an initiator relationship with the drive, and this Unit Attention condition persists for each initiator until the condition is cleared (flag negated) by each initiator individually. The Unit Attention condition results when one of the following events occur:

1. A power-on sequence occurs.
2. A reset is generated internally by the drive (caused by a power glitch).
3. A Bus Device Reset message causes the drive to reset itself.
4. The RESET I/O line resets the drive.
5. An initiator changes one or more of the Mode Select parameters in the drive (these changes could affect one or more of the other initiators).
6. The inquiry data has been changed.
7. The mode parameters in effect for an initiator have been restored from nonvolatile memory.
8. An event occurs that requires the attention of the initiator.
9. A Clear Queue message received.
10. The Log parameters are changed. Unit Attention Condition is posted for all initiators in the system other than the one that changed the Log Parameters.

The Unit Attention Parameters page (page 00h, bit 4 of byte 2) of the Mode Select Command controls whether or not a Check Condition Status is to be reported to affected initiators when a Unit Attention condition exists. See Table 5.2.1-31.

The Unit Attention condition for a particular initiator is cleared when that initiator does one of the following:

1. It sends a Request Sense Command.
2. It sends any other legitimate command, with the exception of the Inquiry command. The Inquiry command does not clear the Unit Attention condition.

When a Unit Attention condition flag is stored in the drive for an initiator, the commands that initiator issues to the drive operate as described in the following paragraphs.

If an initiator sends an Inquiry command to the drive when the drive has stored a Unit Attention condition flag for that initiator (before or after the drive reports Check Condition status), the drive shall perform the Inquiry command and shall not clear the Unit Attention condition.

If an initiator sends a Request Sense command to the drive when a Unit Attention condition flag is stored for that initiator (before or after the drive reports Check Condition), the drive shall discard any pending sense data, report the Unit Attention Sense Key, and clear the Unit Attention condition (negate the flag) for that initiator.

If an initiator issues a command other than Inquiry or Request Sense while a Unit Attention condition flag is stored for that initiator, the drive may or may not perform the command and report Check Condition status, depending on whether or not the Unit Attention bit is zero or one in the Unit Attention Mode Parameters page (Page 00h, bit 4 of byte 2). See Table 5.2.1-31. If a Request Sense is issued next, the Unit Attention condition is reported and cleared (flag negated) as noted in the preceding paragraph. If another command other than Request Sense or Inquiry is issued instead, the drive shall perform the command and return the appropriate status. The Unit Attention condition for the subject initiator is cleared (flag negated) and the sense data and flag indicating there has been a Unit Attention condition are lost.

4.7 Queued I/O processes

Queuing of I/O processes allows a drive to accept multiple commands for execution at a later time.

There are two methods for implementation of queuing, tagged and untagged. Tagged queuing allows the drive to accept multiple commands from each initiator. Untagged queuing allows the drive to accept one command from each initiator. Drives that have SCSI-2 implementation support tagged queuing while in SCSI-1 or SCSI-2 mode. They can use untagged queuing mode if the initiator does not send queue tag messages.

Initiators may add or delete commands to the queue for the a drive within the limitations specified in this specification. When adding a command, the initiator may specify fixed order of execution, allow the drive to define the order of execution, or specify that the command is to be executed next. See glossary in Section 2.0 for terminology definitions when reading the following explanations.

4.7.1 Untagged queuing

Untagged queuing allows the drive to accept a command from an initiator while a command from another initiator is being executed. Only one command for each I T L nexus may be accepted at a time.

A new I/O process may be initiated any time the BUS FREE phase exists even if another I/O process from a different initiator is being executed. If the disconnect privilege is not granted, the drive returns BUSY status to the new I/O process.

The I T L nexus specifies the relationship so that the drive can always reconnect to the initiator to restore the pointers for I/O process as long as only one command per I T L nexus is issued. It is the responsibility of the initiator to assure that only one command is issued at any time.

4.7.2 Tagged queuing

Tagged queuing allows a drive to accept multiple commands from the same or different initiators until the drive's I/O process queue is full. A new I/O process may be initiated any time the BUS FREE phase exists, if the disconnect privilege is granted. If the disconnect privilege is not granted for a tagged command the drive returns BUSY status to the new I/O process.

The queue tag messages (see Table 3.5.3-2) allow the initiator to establish a unique I T L Q nexus to identify each I/O process. Each I/O process may be a command or a set of linked commands with a unique queue tag.

The I T L Q nexus allows the target to reconnect to the desired I/O process and the initiator to restore the correct set of pointers. An initiator may have several I/O processes ongoing to the same or different logical unit as long as each has a unique nexus.

If only SIMPLE QUEUE TAG messages are used, the drive may execute the commands in any order that is deemed desirable within the constraints of the queue management algorithm specified in the control mode page (see Table 5.2.1-28). The command ordering is done by the drive to meet its performance and functional goals. The algorithm used by the drive attempts to achieve certain drive or system performance goals established in the drive firmware for the queued commands and guarantee that all commands will be executed. One possible goal would be to minimize seek times, but there could be others, possibly designed to meet some special system need. Commands from other initiators are also executed in an order selected in the same manner. The drive uses the simple Queue Tag when reconnecting to the initiator.

If ORDERED QUEUE TAG messages are used, the drive executes the commands in the order received with respect to other commands received with ORDERED QUEUE TAG messages. All commands received with a SIMPLE QUEUE TAG message prior to a command received with an ORDERED QUEUE TAG message, regardless of initiator, are executed before that command with the ORDERED QUEUE TAG message. All commands received with a SIMPLE QUEUE TAG message after a command received with an ORDERED QUEUE TAG message, regardless of initiator, are executed after that command with the ORDERED QUEUE TAG message.

A command received with a HEAD OF QUEUE TAG message is placed first in the queue, to be executed next. A command received with a HEAD OF QUEUE TAG message does not suspend an I/O process for which the drive has begun execution. Consecutive commands received with HEAD OF QUEUE TAG messages are executed in a last-in-first-out order.

The control mode page specifies alternative queue management algorithms with additional rules on the order of execution of commands (see 5.2.1-28).

An I/O process received from an initiator without a queue tag message while there are any tagged I/O commands in the command queue from that initiator is an incorrect initiator connection, unless there is a contingent allegiance condition. An I/O process received from an initiator with a queue tag message while there is an untagged command in the command queue from that initiator is also an incorrect initiator connection. In either of these cases the drive removes all commands in the queue from that initiator, aborts the command in process if it is from that initiator, and sets the Sense Key to Aborted Command and the Sense Code to Overlapped Commands Attempted.

The RESERVE and RELEASE commands should be sent with an ORDERED QUEUE TAG message. Use of the HEAD OF QUEUE TAG message with these commands could result in reservation conflicts with previously issued commands.

The TEST UNIT READY and INQUIRY commands are often sent with a HEAD OF QUEUE TAG message, since the information returned is either available or has no effect on the condition of the drive.

The drive recovery option, is to continue execution of commands in the queue after the contingent allegiance condition has cleared. The drive returns BUSY status to all other initiators while the contingent allegiance condition exists. During this time all commands in the queue are suspended. All commands used for recovery operations are untagged commands.

Deferred errors are normally related to a command that has already completed. As such, there is no attempt to return the queue tag value assigned to the original command.

4.8 Parameter rounding

Certain parameters sent to a target with various commands contain a range of values. Targets may choose to implement only selected values from this range. When the target receives a value that it does not support, it either rejects the command (CHECK CONDITION status with ILLEGAL REQUEST sense key) or it rounds the value received to a supported value. The target shall reject unsupported values unless rounding is permitted in the description of the parameter.

Rounding of parameter values, when permitted (Rounding is enabled by Mode Select command, page code 00h, byte 2, bit 2) shall be performed as follows:

A target that receives a parameter value that is not an exact supported value shall adjust the value to one that it supports and shall return CHECK CONDITION status with a sense key of RECOVERED ERROR. The additional sense code shall be set to ROUNDED PARAMETER. The initiator is responsible to issue an appropriate command to learn what value the target has selected.

IMPLEMENTORS NOTE: Generally, the target should adjust maximum-value fields down to the next lower supported value than the one specified by the initiator. Minimum-value fields should be rounded up to the next higher supported value than the one specified by the initiator. In some cases, the type of rounding (up or down) is explicitly specified in the description of the parameter.

4.9 Command Processing Considerations and Exception conditions

The following sections describe some exception conditions and errors associated with command processing and the sequencing of commands.

4.9.1 Programmable Operating definition

Some applications require that the operating definition of a logical unit be modified to meet the special requirements of a particular initiator. The program-controlled modification of the operating definition is provided to allow operating systems to change the operating definition of a more recently developed targets to one which is more compatible with the operating system. This ability requires that the system comply with the low-level hardware definitions of SCSI-2.

The parameters that can be changed by modifying the operating definition of a logical unit include the vendor identification, the device type, the device model, the SCSI compliance level, the SCSI specification level, the command set, and other parameters. The low-level hardware parameters including signal timing and parity definitions cannot be changed by modifying the operating definition. The present operating definition of a logical unit with respect to an initiator can be determined at any time by execution of an Inquiry command. In some vendor-specific cases, it may also be necessary to perform other commands including Mode Sense and Read Capacity.

Each logical unit begins at a particular operating definition. If the logical unit supports the Change Definition command, the present operating definition can be changed to any other operating definition supported by the logical unit. The actual details of the operating definition of a logical unit are vendor-specific. If the operating definition is changed to one that does not include the Change Definition command, the target continues to accept the Change Definition command.

If an error occurs during execution of a Change Definition command, the original operating definition remains in effect after the command is executed. The new operating definition becomes active only after successful execution of the Change Definition command.

Since new operating definitions may preclude the execution of I/O processes that are already in progress, the target may disconnect to allow completion of any I/O processes that are in progress. Operating definition changes that may cause conflicts with the normal operation from other initiators shall be indicated to those initiators by generating a unit attention condition for each other initiator. The additional sense code shall be set to CHANGED OPERATING DEFINITION.

An initiator may request a list of the operating definitions that the target supports and descriptive text for each operating definition using the INQUIRY command.

4.9.2 Incorrect initiator connection

An incorrect initiator connection occurs on a reconnection if:

- (1) an initiator attempts to reconnect to an I/O process, and
- (2) a soft reset condition has not occurred, and
- (3) the initiator does not send an ABORT, ABORT TAG, BUS DEVICE RESET, CLEAR QUEUE, CONTINUE I/O PROCESS, or TERMINATE I/O PROCESS message during the same MESSAGE OUT phase as the IDENTIFY message.

An incorrect initiator connection also occurs on an initial connection when an initiator:

- (1) attempts to establish an I_T_L_Q nexus when an I_T_L nexus already exists from a previous connection, or
- (2) attempts to establish an I_T_L nexus when an I_T_L_Q nexus already exists unless there is a contingent allegiance or extended contingent allegiance condition present for the logical unit or target routine.

A target that detects an incorrect initiator connection shall abort all I/O processes for the initiator on the logical unit or target routine and shall return CHECK CONDITION status. The sense key shall be set to ABORTED COMMAND and the additional sense code shall be set to OVERLAPPED COMMANDS ATTEMPTED.

An incorrect initiator connection may be indicative of a serious error and, if not detected, could result in an I/O process operating with a wrong set of pointers. This is considered a catastrophic failure on the part of the initiator. Therefore, host-specific error recovery procedures may be required to guarantee the data integrity on the medium. The target may return additional sense data to aid in this error recovery procedure. Also, some targets may not detect an incorrect initiator connection until after the command descriptor block has been received.

4.9.3 Selection of an Invalid Logical unit

The target's response to selection of a logical unit that is not valid is as follows.

In response to an INQUIRY command the target shall return the INQUIRY data with the peripheral qualifier set to the value required in Table 5.1.1-8. In response to any other command except REQUEST SENSE the target shall terminate the command with CHECK CONDITION status. In response to a REQUEST SENSE command the target shall return sense data. The sense key shall be set to ILLEGAL REQUEST and the additional sense code shall be set to LOGICAL UNIT NOT SUPPORTED.

5.0 Command descriptions

Two types of commands are supported by the drive: commands for all devices; and commands for direct access devices. In each of these categories the drive supports only Group 0, Group 1 and Group 2 commands.

5.1 Command descriptions for all device types

5.1.1 Group 0 commands for All Device Types

The drive implements the following Group 0 commands that are applicable for all device types. See Table 5.1.1-1.

Table 5.1.1-1. Group 0 Commands for all device types

Command OP Code*	Command Name	Applicable Section
00h	Test Unit Ready	5.1.1.1
03h	Request Sense	5.1.1.2
12h	Inquiry	5.1.1.3
1Ch	Receive Diag. Results	5.1.1.5
1Dh	Send Diagnostic	5.1.1.6

*Byte 0 of Command Descriptor Block.

5.1.1.1 Test Unit Ready (00h)

The Test Unit Ready command provides a means to verify the logical unit is ready. This is not a request for a self test. If the logical unit (drive) can accept an appropriate medium access command without returning check condition status, the drive returns a Good status. See Table 5.1.1-2 for proper format.

Table 5.1.1-2. Test Unit Ready Command

Bit Byte(s)	7	6	5	4	3	2	1	0
0	0	0	0	0	0	0	0	0
1	Logical Unit No.			0	0	0	0	0
2	0	0	0	0	0	0	0	0
3	0	0	0	0	0	0	0	0
4	0	0	0	0	0	0	0	0
5	0	0	0	0	0	0	Flag	Link [2]

Notes.

[1] Logical Unit No. must be zero. [2] See "Control Bytes", paragraph 4.2.6.

If the drive cannot become operational or is in a state such that an initiator action (e.g., START command) is required to make the unit ready, the drive returns CHECK CONDITION status with a Sense Key of NOT READY. One of several possible Additional Sense codes indicates the reason for the NOT READY condition.

5.1.1.2 Request Sense Command (03h)

Table 5.1.1-3. Request Sense command

Bit Byte(s)	7	6	5	4	3	2	1	0
0	0	0	0	0	0	0	1	1
1	Logical Unit No							
	0	0	0	0	0	0	0	0
2	0	0	0	0	0	0	0	0
3	0	0	0	0	0	0	0	0
4	Allocation Length in Bytes							
5	0	0	0	0	0	0	Flag	Link[1]

The Request Sense command (TABLE 5.1.1-3) requests that the drive transfer sense data to the initiator in the format shown in Table 5.1.1-4. The sense data shall be valid for a Check Condition status returned on the prior command. This sense data shall be preserved by the drive for the initiator until retrieval is requested by the initiator when it sends the Request Sense command or until the receipt of any other command for the same drive from the initiator that issued the command that resulted in the Check Condition status. Sense data shall be cleared upon receipt of any subsequent command to the drive from the initiator receiving the Check Condition status.

If a drive sends a Check Condition status as a response to a Request Sense command being in error, it shall do so only if the error was a fatal error. For example:

1. The drive receives a nonzero reserved bit in the command descriptor block.
2. An unrecovered parity error occurs on the Data Bus.
3. A drive malfunction prevents return of sense data.

Notes. [1] See "Control Bytes", paragraph 4.2.6.

If any nonfatal error occurs during execution of Request Sense, the drive shall return sense data with Good status. Following a fatal error on a Request Sense command, sense data may be invalid.

The Allocation Length in byte four of the format shown specifies the number of bytes the initiator has allocated for returned sense data. The Allocation Length should always be at least 18 bytes for drive devices for the initiator to receive all of the drive sense data. Any other value indicates the maximum number of bytes that shall be transferred. The drive shall terminate the Data In phase when allocation length bytes have been transferred or when all available sense data have been transferred to the initiator, whichever is less. The drive always returns sense data in the Extended Sense Data Format.

Extended Sense Data Format

The drive is capable of sending 18 bytes of extended sense data, and does send 18 bytes if the Allocation Length of the Request Sense Command is equal to or greater than 18 bytes (otherwise, the number of bytes specified by the Allocation Length are sent). The Extended Sense Data Format is summarized in Table 5.1.1-4. Numbers in parentheses in the sense data are references to notes following the table. A “1” or a “0” means the data bit is always logic 1 or logic 0, respectively, when sent by a drive.

Table 5.1.1-4. Drive Extended Sense Data Summary

Bit Byte(s)	7	6	5	4	3	2	1	0
0	Validity Bit [1]	1 1	1 1	1 1	0 0	0 0	0 0	0 OR 1 [2]
	Error Code							
1	0	0	0	0	0	0	0	0 [3]
	Segment Number							
2	0 Filemark [4]	0 EOM [5]	0 ILI [6]	0	Sense Key			[7]
3	Information Byte (MSB)							[8]
4	:	:						
5	:	:						
6	Information Byte (LSB)							[8]
7	Additional Sense Length - Decimal 10 (Max)							[9]
8	Bytes 8 - 11 are for Command Specific Data.							
9								
10								
11								
12	Additional Sense Code							[10]
13	Additional Sense Code Qualifier							[10]
14	Reserved for Seagate internal use only							[11]
15	SKSV [12]	Sense Key Specific					[13]	
16								
17								
18—n	Product Unique Sense Data							[14]

[] For notes see next page.

continued from previous page

Notes:

- [1] Validity Bit - "1" if the Information Bytes (Bytes 3-6) are valid, "0" if not valid.
- [2] Error Code 70h means current error. Code 71h means a deferred error. See paragraph 5.1.1.2.1 for a description of these two error conditions.
- [3] Segment Number - Always Zeros
- [4] Filemark - Always "0" for drive.
- [5] EOM - End of medium indicator. Always "0" for drive.
- [6] ILI - Incorrect Length Indicator. The requested (previous command) block of data did not match the logical block length of the data on the medium.
- [7] Sense Key - Indicates the general error category. These are listed in Table 5.1.1-5. the code given in byte 12 provides additional clarification of errors. See also note [10] below for related information.
- [8] If the validity bit is a 1, the Information Bytes contain the unsigned Logical Block Address associated with the Sense Key. Unless otherwise specified, the Information Bytes contain the address of the current logical block. For example, if the Sense Key is Medium Error, it is the Logical Block Address of the failure block.
- [9] Additional Length - Specifies additional sense bytes are to follow. This is limited to a maximum of 10 (decimal) additional bytes. If the Allocation length of the Command Descriptor Block is too small to transfer all of the additional sense bytes, the additional sense length is not adjusted to reflect the truncation.
- [10] Additional Sense Code and Additional Sense Code Qualifier - Provides additional clarification of errors whenever Sense Key is valid. Error code definitions are in Table 5.1.1-6. If the condition is not reportable by the drive, the additional Sense Code and Additional Sense Code Qualifier are set to "No Additional Sense Information (Code 0000)".
- [11] Defined for Seagate internal use only.
- [12] [13] The SKSV bits and sense-key specific bytes are described below.

The additional sense bytes field may contain command specific data, peripheral device specific data, or vendor-specific data that further defines the nature of the CHECK CONDITION status.

The sense-key specific field is defined by this specification when the value of the sense-key specific valid (SKSV) bit is one. The definition of this field is determined by the value of the sense key field. This field is reserved for sense keys not described in Table 5.1.1-5.

If the sense key field is set to ILLEGAL REQUEST (5H) and the SKSV bit is set to one, the sense-key specific field is defined as shown in Table 5.1.1-4a. These fields point to illegal parameters in command descriptor blocks and data parameters sent by the initiator in the DATA OUT phase.

continued from previous page

Notes [12] [13]: (continued)

Table 5.1.1-4a. Field Pointer Bytes

Bit Byte(s)	7	6	5	4	3	2	1	0	
15	SKSV	C/D	Rsvd	Rsvd	BPV	Bit Pointer			
16	(MSB)	Field Pointer							
17								(LSB)	

A command data (C/D) bit of one indicates that the illegal parameter is in the command descriptor block. A C/D bit of zero indicates that the illegal parameter is in the data parameters sent by the initiator during the DATA OUT phase.

A bit pointer valid (BPV) bit of zero indicates that the value in the bit pointer field is not valid. A BPV bit of one indicates that the bit pointer field specifies the bit of the byte designated by the field pointer that is in error. When a multiple-bit field is in error, the bit pointer field shall point to the most-significant (left-most) bit of the field.

The field pointer field indicates the byte of the command descriptor block or of the parameter data that was in error. Bytes are numbered starting from zero, as shown in the tables describing the commands and parameters. When a multiple-byte field is in error, the pointer shall point to the most-significant (left-most) byte of the field.

Note. Bytes identified as being in error are not necessarily the place that has to be changed to correct the problem.

If the sense key is RECOVERED ERROR (1h), HARDWARE ERROR (4h) or MEDIUM ERROR (3h) and if the SKSV bit is one, the sense-key specific field is defined as shown in Table 5.1.1-4b. These fields identify the actual number of retries used in attempting to recover from the error condition.

Table 5.1.1-4b. Actual Retry Count Bytes

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

The actual retry count field returns implementation specific information on the actual number of retries used in attempting to recover an error or exception condition.

Not all drives implement reporting actual retry count in bytes 15, 16 and 17. See specific drive Product Manual Vol. 1.

Note. This field relates to the retry count fields specified within the error recovery parameters page of the MODE SELECT command. See Table 5.2.1-26 and paragraph 6.0.

continued from previous page

Notes [12] [13]: (continued)

If the sense key is NOT READY and the SKSV bit is one, the sense-key specific field shall be defined as shown in Table 5.1.1-4c. These fields are only defined for the FORMAT UNIT command with the Immed bit set to one.

Table 5.1.1-4c. Format Progress Indication Bytes

Bit Byte(s)	7	6	5	4	3	2	1	0
15	SKSV			Reserved				
16	MSB			Progress Indication				
17								LSB

The progress indication field is a percent complete indication in which the returned value is the numerator that has 65536 (10000h) as its denominator. The progress indication shall be based upon the total format operation including any certification or initialization operations.

Support or non-support for format progress indication is given in each drive Product Manual (Volume 1).

Note [14]: Bytes 18-n not presently used.

Table 5.1.1-5 lists the Sense Keys in the extended sense data format that are used by the drive.

Table 5.1.1-5. Applicable Drive Sense Keys

Sense Key	Description
0h	No Sense - Indicates there is no specific Sense Key information to be reported for the drive. This would be the case for a successful command or when the ILI bit is one.
1h	Recovered Error - Indicates the last command completed successfully with some recovery action performed by the drive. When multiple recovered errors occur, the last error that occurred is reported by the additional sense bytes. NOTE: For some Mode settings, the last command may have terminated before completing.
2h	Not Ready - Indicates the logical unit addressed cannot be accessed. Operator intervention may be required to correct this condition.
3h	Medium Error - Indicates the command terminated with a nonrecovered error condition, probably caused by a flaw in the medium or an error in the recorded data.
4h	Hardware Error - Indicates the drive detected a nonrecoverable hardware failure while performing the command or during a self test. This includes SCSI interface parity error, controller failure, device failure, etc.
5h	Illegal Request - Indicates an illegal parameter in the command descriptor block or in the additional parameters supplied as data for some commands (Format Unit, Mode Select, etc). If the drive detects an invalid parameter in the Command Descriptor Block, it shall terminate the command 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 IDENTIFY message was received. This could also indicate an attempt to write past the last logical block.
6h	Unit Attention - Indicates the drive may have been reset. See Paragraph 4.6 for more detailed information about the Unit Attention Condition.
7h	Data Protect - Indicates that a command that reads or writes the medium was attempted on a block that is protected from this operation. The read or write operation is not performed.
9h	Firmware Error - Vendor specific sense key.
Bh	Aborted Command - Indicates the drive aborted the command. The initiator may be able to recover by trying the command again.
Ch	Equal - Indicates a SEARCH DATA command has satisfied an equal comparison.
Dh	Volume Overflow - Indicates a buffered peripheral device has reached the end of medium partition and data remains in the buffer that has not been written to the medium.
Eh	Miscompare - Indicates that the source data did not match the data read from the medium.

Table 5.1.1-6 lists the Extended Sense, Additional Sense and Additional Sense Qualifier Codes.

Note: Table 5.1.1-6 is for reference only, as not all drives listed on the cover of this manual support all of the codes listed.

Table 5.1.1-6. Error codes for bytes 12 and 13 of Sense Data (values are in hexadecimal)

Byte 12	Byte 13	Description
00	00	No Additional Sense Information
01	00	No Index/Logical Block Signal
02	00	No Seek Complete
03	00	Peripheral Device Write Fault
04	00	Logical Unit Not Ready, Cause Not Reportable
04	01	Logical Unit Not Ready, Becoming Ready
04	02	Logical Unit Not Ready, Initializing Command Required
04	03	Logical Unit Not Ready, Manual Intervention Required
04	04	Logical Unit Not Ready, Format in Progress
08	00	Logical Unit Communication Failure
08	01	Logical Unit Communication Time-Out
08	02	Logical Unit Communication Parity Error
09	00	Track Following Error
09	01	Servo Fault
09	04	Head Select Fault
0A	00	Error log overflow
0C	00	Write error
0C	01	Write Error Recovered With Auto-Reallocation
0C	02	Write Error - Auto Reallocation failed
10	00	ID CRC Or ECC Error
11	00	Unrecovered Read Error
11	01	Read Retries Exhausted
11	02	Error Too Long To Correct
11	04	Unrecovered Read Error - Auto Reallocation Failed
12	00	Address Mark Not Found For ID Field
12	01	Recovered Data without ECC using Previous Logical Block ID
12	02	Recovered Data with ECC using Previous Logical Block ID
14	00	Logical Block Not Found
14	01	Record Not Found
15	00	Random Positioning Error
15	01	Mechanical Positioning Error
15	02	Positioning Error Detected By Read Of Medium
16	00	Data Synchronization Mark Error
17	00	Recovered Data With No Error Correction Applied
17	01	Recovered Data Using Retries
17	02	Recovered Data Using Positive Offset
17	03	Recovered Data Using Negative Offset
17	05	Recovered Data Using Previous Logical Block ID
17	06	Recovered Data Without ECC - Data Auto Reallocated
18	00	Recovered Data With ECC
18	01	Recovered Data With ECC And Retries Applied
18	02	Recovered Data With ECC And/Or Retries, Data Auto-Reallocated
18	05	Recovered Data with ECC and/or retries
19	00	Defect List Error
19	01	Defect List Not Available
19	02	Defect List Error In Primary List
19	03	Defect List Error in Grown List
1A	00	Parameter List Length Error
1B	00	Synchronous Data Transfer Error
1C	00	Defect List Not Found
1C	01	Primary Defect List Not Found
1C	02	Grown Defect List Not Found
1C	83	Seagate Unique diagnostic code
1D	00	Miscompare During Verify Operation

continued from previous page

Byte	Byte	Description
12	13	
20	00	Invalid Command Operation Code
21	00	Logical Block Address Out Of Range
24	00	Invalid Field In CDB
25	00	Logical Unit Not Supported
26	00	Invalid Field In Parameter List
26	01	Parameter Not Supported
26	02	Parameter Value Invalid
26	03	Invalid Field Parameter - Threshold Parameter
26	98	Invalid Field Parameter - Check Sum
26	99	Invalid Field Parameter - Firmware Tag
27	00	Write Protected
28	00	Not-Ready to Ready Transition; medium may have changed
29	00	Power On, Reset, Or Bus Device Reset Occurred
29	01	Power-On Reset occurred
2A	00	Mode Parameters Changed
2A	01	Mode Parameters Changed
2A	02	Log Parameters Changed
2C	00	Command Sequence Error
2F	00	Tagged Commands Cleared By Another Initiator
31	00	Medium Format corrupted
31	01	Format Failed
31	91	Format Corrupted, World Wide Name (WWN) is invalid
32	00	No Defect Spare Location Available
32	01	Defect list update Error
32	02	No spares available - Too many defects on one Track
37	00	Parameter Rounded
3D	00	Invalid Bits in Identify Message
3E	00	Logical Unit has Not Self Configured yet
3F	00	Target Operating Conditions have changed
3F	01	Microcode Changed
3F	02	Changed Operating Definition
3F	90	Invalid APM parameters
3F	91	World Wide Name (WWN) mismatch
40	01	DRAM Parity Error
42	00	Power-On or Self-Test Failure
43	00	Message Error
44	00	Internal Target Failure
45	00	Select/Reselection Failure
47	00	SCSI Parity Error
48	00	Initiator Detected Error Message Received
49	00	Invalid Message Error
4C	00	Logical Unit Failed Self-Configuration
4E	00	Overlapped Commands Attempted
5B*	00	Log Exception
5B*	01	Threshold Condition Met
5B*	02	Log Counter At Maximum
5B*	03	Log List Codes EXHAUSTED
5C	00	RPL Status Change
5C	01	Spindles Synchronized
5C	02	Spindles Not Synchronized
5D	00	Failure Prediction Threshold exceeded
65	00	Voltage Fault
80	00	General Firmware Error Qualifier
81	00	Reassign Power - fail recovery failed

***Note.** Can be supported, but is a factory installed option.

5.1.1.2.1 Deferred errors

Error code 70h (current error) indicates that the CHECK CONDITION or COMMAND TERMINATED status returned is the result of an error or exception condition on the command that returned the CHECK CONDITION or COMMAND TERMINATED status or an unexpected bus free condition. This includes errors generated during execution of the command by the actual execution process. It also includes errors not related to any command that are first observed during execution of a command. Examples of this latter type of error include disc servo-mechanism off-track errors and power-up test errors.

Error code 71h (deferred error) indicates that the CHECK CONDITION status returned is the result of an error or exception condition that occurred during execution of a previous command for which GOOD status has already been returned. Such commands are associated with use of the immediate bit (start unit), with some forms of caching, and with multiple command buffering. Targets that implement these features are required to implement deferred error reporting.

When the drive does not use the AEN (Asynchronous Event Notification) Feature, the deferred error may be indicated by returning CHECK CONDITION status to the appropriate initiator as described below. The subsequent execution of a REQUEST SENSE command shall return the deferred error sense information.

If CHECK CONDITION status for a deferred error is returned, the current command has not performed any storage operations or output operations to the media. After the target detects a deferred error condition on a logical unit, it shall return a deferred error according to the rules described below:

(1) If a deferred error can be recovered with no external system intervention, a deferred error indication shall not be posted unless required by the error handling parameters of the MODE SELECT command. The occurrence of the error may be logged if statistical or error logging is supported.

(2) If a deferred error can be associated with a causing initiator and with a particular function or a particular subset of data, and the error is either unrecovered or required to be reported by the mode parameters, a deferred error indication shall be returned to the causing initiator. If an initiator other than the causing initiator attempts access to the particular function or subset of data associated with the deferred error, a BUSY status shall be returned to that initiator in response to the command attempting the access.

Note. Not all devices may be sufficiently sophisticated to identify the function or data that has failed. Those that cannot should treat the error in the following manner.

(3) If a deferred error cannot be associated with a causing initiator or with a particular subset of data, a deferred error indication shall be returned on behalf of the failing logical unit to each initiator. If multiple deferred errors have accumulated for some initiators, only the last error shall be returned.

(4) If a current command has not yet started executing, and a deferred error occurs, the command shall be terminated with CHECK CONDITION status and deferred error information posted in the sense data. By convention, the current command is considered to have started execution if the target has changed phase from the COMMAND phase to the next normal phase of the command sequence. If a deferred error occurs while a current command is executing and the current command has been affected by the error, the command shall be terminated by CHECK CONDITION status and current error information shall be returned in the sense data. In this case, if the current error information does not adequately define the deferred error condition, a deferred error may be returned after the current error information has been recovered. If a deferred error occurs while a current command is executing and the current command completes successfully, the target may choose to return the deferred error information after the completion of the current command.

Deferred errors may indicate that an operation was unsuccessful long after the command performing the data transfer returned GOOD status. If data that cannot be replicated or recovered from other sources is being stored using such buffered write operations, synchronization commands should be performed before the critical data is destroyed in the host initiator. This is necessary to be sure that recovery actions can be taken if deferred errors do occur in the storing of the data. When AEN is not implemented, the synchronizing process provides the necessary commands to allow returning CHECK CONDITION status and subsequent returning of deferred error sense information after all buffered operations are guaranteed to be complete.

5.1.1.3 Inquiry Command (12h)

The INQUIRY command requests that information regarding parameters of the drive be sent to the initiator. An option Enable Vital Product Data (EVPD) allows the initiator to request additional information about the drive. See paragraph 5.1.1.3.1. Several Inquiry commands may be sent to request the vital product data pages instead of the standard data shown in Table 5.1.1-8.

Table 5.1.1-7. Inquiry command (12h)

Bit Byte(s)	7	6	5	4	3	2	1	0
0	0	0	0	1	0	0	1	0
1	Logical Unit No.[1] 0 0 0			0	0	0	0	EVPD [2]
2	Page Code							[3]
3	0	0	0	0	0	0	0	0
4	Allocation Length In Bytes							[4]
5	0	0	0	0	0	0	Flag	Link [5]

Notes.

- [1] LUN must be zero.
- [2] An enable vital product data (EVPD) bit of one specifies that the drive shall return the vital product data specified by the page code field. An EVPD bit of zero specifies that the drive shall return the standard INQUIRY data.
- [3] The page code field specifies the page of vital product data information that the drive shall return. If the EVPD bit is zero and the page code field is not zero the drive shall return CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and an additional sense code of INVALID FIELD IN CDB.
- [4] The Allocation Length specifies the number of bytes that the initiator has allocated for returned data. The drive shall terminate the Data-In phase when Allocation Length bytes have been transferred or when all available data have been transferred to the initiator, whichever is less.
- [5] See "Control Byte" paragraph 4.2.6.

The INQUIRY command returns CHECK CONDITION status only when the drive cannot return the requested INQUIRY data.

If an INQUIRY command is received from an initiator with a pending unit attention condition (i.e., before the drive reports CHECK CONDITION status), the drive performs the INQUIRY command and does not clear the unit attention condition.

The INQUIRY command is typically used by the initiator after a reset or power-up condition to determine the device types for system configuration. To minimize delays after a reset or power-up condition, the standard INQUIRY data is available without incurring any media access delays. Since the drive stores some of the INQUIRY data on the device media it may return zeros or ASCII spaces (20h) in those fields until the data is available.

The INQUIRY data may change as the drive executes its initialization sequence or in response to a CHANGE DEFINITION command. For example, the drive may contain a minimum command set in its nonvolatile memory and load its final firmware from the medium when it becomes ready. After it has loaded the firmware it may support more options and therefore return different supported options information in the INQUIRY data.

Drive Inquiry Data

The drive standard INQUIRY data contains 36 required bytes, followed by a number of bytes of drive specific data that is drive dependent. See individual drive Product Manual. The standard INQUIRY data is given in Table 5.1.1-8.

Table 5.1.1-8. Drive Inquiry Data Format

Bit Byte(s)	7	6	5	4	3	2	1	0
0	Peripheral Qualifier[1]		Peripheral Device Type [1]					
1	0	0	0	0	0	0	0	0
2	0	0	0	0	0	ANSI-Approved Ver.		[2]
3	AENC 0 [3]	TRMIOP 0 [4]	0	0	Response Data Format			[5]
4	Additional Length							[6]
5	Reserved							
6	0	0	0	0	0	0	ADR32 [21]	ADR16 [22]
7	RELADR [7]	WBUS 32 [8]	WBUS 16 [9]	SYNC [10]	LINK- ED [11]	TrnDis [23]	CMD QUE [12]	SOFT RE- SET [13]
8 : 15	Vendor Identification							[14]
16 : 31	Product Identification							[15]
32 : 35	Product Revision Level							[16]
36 : 43	Drive Serial Number							[17]
44 : 55	Unused Vendor Specific Area (00h)							
56 : 95	Reserved (00h)							[18]
96 : 143	Copyright Notice							[19]
144 : 147	Distribution Diskette Serial Number							[20]

Notes.

- [1] The Peripheral Qualifier and Peripheral Device Type field value of 00h indicates a direct-access device (magnetic disc) is connected to this logical unit.
- [2] The ANSI-Approved Version field indicates the implemented version of this standard and is defined below in Table 5.1.1-9.

Table 5.1.1-9. ANSI-approved version

Code	Description
0h	The device might or might not comply to an ANSI approved standard.
1h	The device complies to ANSI X3.131-1986 (SCSI-1).
2h	The device complies to ANSI IX3.131-199x (SCSI-2). This code is reserved to designate this standard upon approval by ANSI.

- [3] An asynchronous event notification capability (AENC) bit of zero indicates that the drive does not support the asynchronous event notification capability.
- [4] A terminate I/O process (TRMIOP) bit of zero indicates that the drive does not support the TERMINATE I/O PROCESS message.
- [5] A Response Data Format value of zero indicates the INQUIRY data format is as specified in the ANSI SCSI-1 standard. A response data format value of one indicates compatibility with some products that were designed prior to the development of the ANSI SCSI-2 standard (i.e. CCS). A response data format value of two indicates that the data shall be in the format specified in the SCSI-2 standard. Response data format values greater than two are reserved.
- [6] The Additional Length field shall specify the length in bytes of the parameters. If the allocation length of the command descriptor block is too small to transfer all of the parameters, the additional length is not adjusted to reflect the truncation.
- [7] A relative addressing (RELADR) bit of one indicates that the drive supports the relative addressing mode. If this bit is set to one the linked command (LINKED) bit shall also be set to one since relative addressing can only be used with linked commands. A RELADR bit of zero indicates the drive does not support relative addressing.
- [8] A wide bus 32 (WBus32) bit of one indicates that the drive supports 32-bit wide data transfers. A value of zero indicates that the drive does not support 32-bit wide data transfers.
- [9] A wide bus 16 (WBus16) bit of one indicates that the drive supports 16-bit wide data transfers. A value of zero indicates that the drive does not support 16-bit wide data transfers.

Note. If the values of both the WBus 16 and WBus 32 bits are zero the drive only supports 8-bit wide data transfers.

- [10] A synchronous transfer (SYNC) bit of one indicates that the drive supports synchronous data transfer. A value of zero indicates the drive does not support synchronous data transfer.
- [11] A linked command (LINKED) bit of one indicates that the drive supports linked commands. A value of zero indicates the drive does not support linked commands.
- [12] A command queuing (CMDQUE) bit of one indicates that the drive supports tagged command queuing. A value of zero indicates the drive does not support tagged command queuing.
- [13] Soft RESET bit of zero indicates that the drive responds to the reset condition with the hard RESET alternative (see section 3.2.2).

A Soft RESET bit of one indicates that the drive responds to the RESET condition with the soft RESET alternative (see section 3.2.2).

continued from previous page

- [14] The Vendor Identification field contains the ASCII data giving vendor name ("SEAGATE").
- [15] The Product Identification field contains ASCII data giving the drive model number. The data shall be left-aligned within this field.
- [16] The Product Revision Level field contains the four bytes of ASCII data "XXXX", where the value of XXXX is the last 4 digits of the Firmware Release number.
- [17] Drive Serial Number field contains the 8 bytes of ASCII data "XXXXXXXX", where the value XXXXXXXX is the drive serial number.
- [18] The reserved area from byte 56 through byte 95 is filled with 00h.
- [19] The Copyright Notice field contains the 48 bytes of ASCII data "Copyright (c) 199X Seagate All rights reserved", where "X" indicates the current year.
- [20] The Distribution Diskette Serial Number field contains the 4 bytes of ASCII data "XXXX", where the value XXXX is the Distribution Diskette serial number.
- [21] A one bit indicates the drive supports 32-bit wide SCSI addresses. Zero indicates non-support.
- [22] A one bit indicates the drive supports 16-bit wide SCSI addresses. Zero indicates non-support. A zero in both bits [21] and [22] indicates the drive supports only 8-bit SCSI addresses.
- [23] TrnDis. Transfer Disable bit. When this bit is one it indicates that the drive supports the Continue I/O Process and Target Transfer Disable messages. When this bit is zero it indicates the drive does not support the above two messages.

5.1.1.3.1 Vital product data pages

The initiator requests the vital product data information by setting the EVPD bit to one and specifying the page code of the desired vital product data. If the drive does not implement the requested page it shall return CHECK CONDITION status. The sense key shall be set to ILLEGAL REQUEST and the additional sense code shall be set to INVALID FIELD IN CDB.

This section describes the vital product data page structure and the vital product data pages that are applicable to the drive. These pages are optionally returned by the INQUIRY command and contain specific product information about the drive. The vital product data includes unit serial numbers, device operating definitions, firmware release numbers, servo ROM and RAM release numbers and the date code from the manufacturer's defect log.

Table 5.1.1-10a. Supported Vital Product Data Pages (0h)

Bit Byte(s)	7	6	5	4	3	2	1	0
0	Peripheral Qualifier [1]			Peripheral Device Type				[1]
1	Page Code (00h)							[2]
2	0	0	0	0	0	0	0	0
3	Page Length							[3]
4 : N	Supported Page List							[4]

Notes.

- [1] The Peripheral Qualifier field and the Peripheral Device Type field are as previously defined.
- [2] Page 0h provides a list of all supported vital product data pages. The Page Code field shall be set to the value of the page code field in the INQUIRY command descriptor block.
- [3] The Page Length field specifies the length of the supported page list. If the allocation length is too small to transfer all of the page, the page length shall not be adjusted to reflect the truncation. The page length reported by ASA-1 firmware is 06. The page length reported by ASA-2 firmware is 07. The total number of bytes returned (N) will be Page Length +4 in either case.

continued from previous page

- [4] The Supported Page List field shall contain a list of all vital product data page codes implemented for the drive in ascending order beginning with page code 0h. See Table 5.1.1-10b.

Table 5.1.1-10b. Vital Product Data Page Codes

Page Code	Description
00h	Supported Vital Product Data Pages
80h	Unit Serial Number Page
81h	Implemented Operating Definitions Page
C0h	Firmware Numbers Page
C1h	Date Code Page
C2h	Jumper Settings Page
C3h	Device Behavior Page (only returned by devices with ASA-2 firmware)

Table 5.1.1-10c. Unit Serial Number Page (80h)

Bit Byte(s)	7	6	5	4	3	2	1	0
0	Peripheral Qualifier [1]			Peripheral Device Type				[1]
1	Page Code (80h)							[2]
2	0	0	0	0	0	0	0	0
3	Page Length (14h)							[3]
4 : 11	Product Serial Number							[4]
12 : 23	Board Serial Number							[5]

Notes.

- [1] The Peripheral Qualifier field and the Peripheral Device Type field are as defined previously.
- [2] Page 80h provides the product serial number and product circuit board number for the drive (See Table 5.1.1-10c).
- [3] The Page Length field specifies the length of the product serial number. If the allocation length is too small to transfer all of the page, the page length shall not be adjusted to reflect the truncation.
- [4] The Product Serial Number field contains ASCII data. The least significant ASCII character of the serial number shall appear as the last byte of this field. If the product serial number is not available, the drive returns ASCII spaces (20h) in this field.
- [5] The board serial number field contains ASCII data that is vendor specific. The least significant ASCII character of the serial number shall appear as the last byte of this field. If the board serial number is not available, the target shall return ASCII spaces (20h) in this field.

Table 5.1.1-10d. Implemented Operating Definition Page (81h)

Bit Byte(s)	7	6	5	4	3	2	1	0
0	Peripheral Qualifier [1]			Peripheral Device Type				[1]
1	Page Code (81h)							[2]
2	0	0	0	0	0	0	0	0
3	Page Length (04h)							[3]
4	0	Current Operating_Definition						[4]
5	SAVIMP	Default Operating Definition						[5]
6	SAVIMP	Supported Operating Definition						[5]
7	SAVIMP	Supported Operating Definition						[5]

Notes.

- [1] The Peripheral Qualifier Field and the Peripheral Device Type field are as defined previously. (See Table 5.1.1-8 notes).
- [2] Page 81h defines the current operating definition, the default operating definition, and which operating definitions are implemented by the drive. These operating definition values are specified in the CHANGE DEFINITION command (see Table 5.1.3-1).
- [3] The Page Length field specifies the length of the following operating definitions. If the allocation length of the command descriptor block is too small to transfer all of the page, the page length shall not be adjusted to reflect the truncation.
- [4] The current operating definition field returns the value of the present operating definition. The default operating definition field returns the value of the operating definition the drive uses when power is applied if no operating definition is saved (see Table 5.1.1-10e).
- [5] For each operating definition there is an associated save implemented (SAVIMP) bit. A SAVIMP bit of zero indicates that the corresponding operating definition parameter cannot be saved. A SAVIMP bit of one indicates that the corresponding operating definition parameter can be saved. The drive saves the default operating definition and all supported operating definitions.

Table 5.1.1-10e. Operating Definition Field

Code	Operating description
00h	Use Current Operating Definition
01h	SCSI X3.131-1986 Operating Definition
03h	SCSI-2 X3.131-199x Operating Definition

Table 5.1.1-10fa. Firmware Numbers Page (C0h)
(Applies to model families ST11200 and ST3600)

Bit Byte(s)	7	6	5	4	3	2	1	0
0	Peripheral Qualifier [1]		Peripheral Device Type					[1]
1	Page Code (C0h)							[2]
2	0	0	0	0	0	0	0	0
3	Page Length (10h)							[3]
4 : 11	SCSI Firmware Release number							[4]
12 : 19	ASCII space characters							
20 : 27	Servo ROM Release number							[4]

Notes.

- [1] The Peripheral Qualifier field and the Peripheral Device Type field are as defined previously.
(See Table 5.1.1-8 notes).
- [2] Page C0h provides the firmware release numbers for the drive (see Table 5.1.1-10fb).
- [3] The Page Length field specifies the length of the product firmware numbers. If the allocation length is too small to transfer all of the page, the page length shall not be adjusted to reflect the truncation.
- [4] The firmware release numbers fields contain ASCII data. The least significant ASCII character of the Drive firmware number shall appear as the last byte of a successful data transfer.

Table 5.1.1-10fb. Firmware Numbers Page (C0h)
(Applies to model families other than those covered by
Table 5.1.1-10fa)

Bit Byte(s)	7	6	5	4	3	2	1	0	
0	Peripheral Qualifier [1]		Peripheral Device Type						[1]
1	Page Code (C0h)								
2	0	0	0	0	0	0	0	0	
3	Page Length (20h)							[3]	
4 : 11	SCSI Firmware Release number							[4]	
12 : 19	Servo RAM Release number							[4]	
20 : 27	Servo ROM Release number							[4]	
28 : 31	Servo RAM Release date in ASCII								
32 : 35	Servo ROM Release date in ASCII								

Notes.

- [1] The Peripheral Qualifier field and the Peripheral Device Type field are as defined previously.
(See Table 5.1.1-8 notes).
- [2] Page C0h provides the firmware release numbers for the drive (see Table 5.1.1-10fb).
- [3] The Page Length field specifies the length of the product firmware numbers. If the allocation length is too small to transfer all of the page, the page length shall not be adjusted to reflect the truncation.
- [4] The firmware release numbers fields contain ASCII data. The least significant ASCII character of the Drive firmware number shall appear as the last byte of a successful data transfer.

Table 5.1.1-10g. ETF Log Date Code Page (C1h)

Bit Byte(s)	7	6	5	4	3	2	1	0
0	Peripheral Qualifier [1]			Peripheral Device Type				[1]
1	Page Code (C1h)							[2]
2	0	0	0	0	0	0	0	0
3	Page Length (06h)							[3]
4 : 9	(MSB) ETF Log Date in ASCII							[4] (LSB)
10 : : 15	(MSB) Compile Date Code							[5] (LSB)

Notes.

[1] The Peripheral Qualifier field and the Peripheral Device Type field are as defined previously. (see Table 5.1.1-8 notes).

[2] Page C1h provides the date code from the drive defect list (see Table 5.1.1-10g).

[3] The Page Length field specifies the length of the product date code. If the allocation length is too small to transfer all of the page, the page length shall not be adjusted to reflect the truncation.

[4] The ETF Log date code field contains ASCII data.

The data is stored in the format MMDDYY. Leading ASCII zero characters are added to single-digit months or days.

[5] Contains 6 ASCII bytes of data for a date of the form MMDDYY.

Table 5.1.1-10h. Jumper Settings Page (C2h)

Bit Byte(s)	7	6	5	4	3	2	1	0
0	Peripheral Qualifier [1]			Peripheral Device Type [1]				
1	Page Code (C2h) [2]							
2	0	0	0	0	0	0	0	0
3	Page Length (02h) [3]							
4	DS [4]	MS [5]	WP [6]	PE [7]	DRIVE ID [8]			
5	Rsvd	Rsvd	Rsvd	Rsvd	Rsvd	Rsvd	Rsvd	Term Enable [9] [10]

Notes.

- [1] The Peripheral Qualifier and Peripheral Device Type are as defined previously. (See Table 5.1.1-8 notes).
- [2] Page C2h provides all the jumper settings for the drive. The bits in byte 4 indicate which jumpers are on.
- [3] The Page Length field specifies the length of the Jumper Setting Page. If the allocation length is too small to transfer all of the page, the page length shall not be adjusted to reflect the truncation.
- [4] Delayed Motor Start (DS) bit when set to 1 indicates that this jumper is on.
- [5] Motor Start (MS) bit when set to 1 indicates that the jumper is on.
- [6] Write Protected (WP) bit when set to 1 indicates that the write protect jumper is on.
- [7] On some drives, Parity Enable (PE) bit when set to 1 indicates that SCSI parity error checking jumper is on. Other drives have a PD (Parity Disable) jumper, where Parity is enabled when the jumper is off. See the drive Product Manual or Installation Guide for information on how the drive of interest is set up.

continued from previous page

[8] The Drive ID is shown below in the table. Bit 3 is the most significant bit and bit 0 is the least significant bit.

Bit 3	Bit 2	Bit 1	Bit 0	Drive ID
0	0	0	0	0
0	0	0	1	1
0	0	1	0	2
0	0	1	1	3
0	1	0	0	4
0	1	0	1	5
0	1	1	0	6
0	1	1	1	7
1	0	0	0	8
1	0	0	1	9
1	0	1	0	10
1	0	1	1	11
1	1	0	0	12
1	1	0	1	13
1	1	1	0	14
1	1	1	1	15

[9] Terminator Enable (TE). Not used on all drives. See Drive Product Manual Vol. 1 for applicability.

[10] It is not presently possible to return information on the Terminator Power (TP) jumpers.

The Device Behavior page (VPD page C3) will be used by the regression tests to determine what behavior should be expected from a particular firmware package.

Table 5.1.1-10i. Device Behavior Page (C3h)

Bit Byte(s)	7	6	5	4	3	2	1	0
00	Peripheral Qualifier [1]			Peripheral Device Type [2]				
01	Page Code (C3h)							
02	Reserved							
03	Page Length (32)							[3]
04	Version Number							[4]
05	Behavior Code							[5]
06	Behavior Code Version Number							[6]
07 .. 22	ASCII Family Number (16 bytes)							[7]
23	Maximum Interleave							[8]
24	Default Number of Cache Segments							[9]
25+	Feature Flags and Additional Byte Fields will go here but are undefined at this time.							

Notes:

[1] The Peripheral Qualifier and the Peripheral Device Type fields are defined in Table 5.1.1-8 Notes.

[2] See Table 5.1.1-8 Notes.

[3] Page Length defines the length of the Device Behavior information in bytes. If the allocation length is too small to transfer all of the page, the page length shall not be adjusted to reflect the truncation.

[4] The Version Number is a one byte short form notation for the 24 byte assignment in the Firmware Numbers page. Version Numbers are registered by Engineering services.

[5] The Behavior Code and Behavior Code Version are jointly assigned by the Firmware Engineering Man-

[6] agers of all SCSI design locations.

[7] The ASCII Model Number is identical to the Product Identification Number given in the Standard Inquiry command data (see Table 5.12.1-8 Notes).

[8] The Maximum Interleave byte specifies the maximum value which the drive can support in the least significant byte of Interleave in the Format Unit command (see Table 5.2.1-3). The actual interleave which has been used during the last Format, providing it has not been changed by an intervening Mode Select command is reported in the Interleave Bytes of the Format Device Bytes of the Format Device Page (see Table 5.2.1-24).

[9] Default Number of Cache Segments is identical to the same parameter given in the Mode Caching page (see Table 5.2.1-27).

5.1.1.4 Copy command (18h)

Not Implemented. If received the drive sends a “Check Condition” status and a Sense Key of “Illegal Request”.

5.1.1.5 Receive Diagnostic Results Command (1Ch)

Table 5.1.1-11. Receive Diagnostic Results Command

Bit Byte(s)	7	6	5	4	3	2	1	0
0	0	0	0	1	1	1	0	0
1	Logical Unit No.[1] 0 0 0			0	0	0	0	0
2	0	0	0	0	0	0	0	0
3	Allocation Length In Bytes (MSB)							[2]
4	Allocation Length In Bytes (LSB)							[2]
5	0	0	0	0	0	0	Flag	Link [3]

The Receive Diagnostic Results command requests analysis data be sent to the initiator after completion of a Send Diagnostic command. The drive supports the optional Page format, wherein the initiator sends additional pages after a Send Diagnostic command. These additional pages have a page code that specifies to the drive the format of the data to be returned after it receives a Receive Diagnostic Results command.

If no data in the optional Page format was requested by the Send Diagnostics command (0 in bit 4 of Table 5.1.1-15), the data returned to the initiator is in the format shown in Table 5.1.1-12.

If the Send Diagnostics command requested either page 00h or page 40h (the only two optional pages supported by the drive), data returned is in the format shown in Table 5.1.1-13 or 5.1.1-14, respectively.

All FRU and error code definitions are unique to this product and intended for Factory/Field Maintenance personnel.

Notes.

[1] LUN must be zero.

[2] The Allocation Length shall specify the number of bytes the initiator has allocated for returned data. An Allocation Length of zero indicates that no data shall be transferred. Any other value indicates the maximum number of bytes that shall be transferred. The drive terminates the Data In phase when Allocation Length bytes have been transferred or when all available data has been transferred to the initiator, whichever is less.

[3] See “Control Byte” paragraph 4.2.6.

Table 5.1.1-12. Diagnostic Data Bytes

Code	Byte	Description	
00h	0	Additional Length (MSB)	[1]
28h	1	Additional Length (LSB)	[1]
XXh	2	FRU Code (most probable)	[2]
XXh	3	FRU Code	[2]
XXh	4	FRU Code	[2]
XXh	5	FRU Code (least probable)	[2]
XXh	6	Error Code (MSB)	[3]
V.U.	7	Error Code (LSB)	[4]
V.U.	8 thru n	Additional Vendor Unique Fault Information	[5]

Notes.**[1] Additional Length:**

This two byte value indicates the number of additional bytes included in the diagnostic data list. For example, if no product unique byte (byte 7) is available, this value would be 0006h. A value of 0000h means that there are no additional bytes.

[2] FRU Code:

A Field Replaceable Unit code is a byte that identifies an assembly that may have failed. The codes will be listed in probability order, with the most probable assembly listed first and the least probable listed last. A code of 00h indicates there is no FRU information and a code of 01h indicates the entire unit should be replaced. Seagate drives return 00h in these bytes.

[3] Error Code:

This two byte value provides information designating which part of a diagnostic operation has failed. The byte 7 error code is vendor unique and defined in note [4]. Usually Seagate drives support only some subset of the list given in note [4] following.

continued from previous page

[4] Vendor Unique

Diagnostic Fault Codes

01h	Formatter Diagnostic Error
02h	Microprocessor RAM Diagnostic Error
04h	No Drive Ready
08h	No Sector or Index Detected
09h	Fatal Hardware Error While Doing Drive Diagnostics
0Ch	No Drive Command Complete
10h	Unable to Set Drive Sector Size
14h	Unable to Clear Drive Attention
18h	Unable to Start Spindle Motor
20h	Unable to Recal Drive
30h	Unable to Send Write Current Data to Drive
34h	Unable to Issue Drive Seek Command
40h	Unable to Read User Table From Drive
41h	Ran Out of Sectors While Doing Drive Diagnostics
42h	Unable to Read Reallocation Table
43h	Unable to Read ETF Log
45h	Firmware Read from Disc or Sent by Host has an Invalid Checksum
60h	Thermal Calibration Failure
70h	Microprocessor Internal Timer Error
80h	Buffer Controller Diagnostic Error
81h	Buffer RAM Diagnostic Error
C1h	Data Miscompare While Doing Drive Diagnostics

[5] Additional Vendor Unique codes (Not Available).

If the Send Diagnostics command requested the Supported Diagnostics Pages list (PF bit = 1), the drive returns data in the format shown in Table 5.1.1-13 after receiving the Receive Diagnostics Results command. It lists all of the diagnostics pages supported by the drive.

Table 5.1.1-13. Supported Diagnostic Pages

Bit Byte(s)	7	6	5	4	3	2	1	0	
0	Page Code (00h)								
1	Reserved								
2	(MSB)	Page Length (n-3)				[1]			
3								(LSB)	
4 : n	Supported Page List				[2]				

Notes.

[1] The page length field specifies the length in bytes of the following supported page list.

[2] The supported page list field shall contain a list of all diagnostic page codes implemented by the drive in ascending order beginning with page code 00h. The drive presently supports only pages 00h (Table 5.1.1-13) and 40h (Table 5.1.1-14).

Translate Address Page

The translate address page allows the initiator to translate a logical block address into a physical sector address or a physical sector address to a logical block address. The address to be translated is passed to the target during the data-out phase associated with the Send Diagnostic command and the results are returned to the initiator during the data-in phase following the Receive Diagnostic Results command. The translated address is returned in the translate address page-Receive Diagnostic (Table 5.1.1-14).

Table 5.1.1-14. Translate Address Page - Receive Diagnostic

Bit Byte(s)	7	6	5	4	3	2	1	0	
0	Page Code (40h)				[1]				
1	Reserved								
2	(MSB)	Page Length				[2]			
3								(LSB)	
4	Reserved				Supplied Format		[3]		
5	RAREA [4]	ALTSEC [5]	ALTTK [6]	RSVD	Translated Format		[7]		
6 : 13	Translated Address				[8]	(if available)			

Notes.

- [1] The translate address page contains a four byte page header which specifies the page code and length followed by two bytes which describe the translated address followed by the translated address.
- [2] The Page Length field contains the number of parameter bytes which follow.
- [3] The Supplied Format field contains the value from the Send Diagnostic command supplied format field (see Table 5.1.1-17).
- [4] A reserved area (RAREA) bit of one indicates that all or part of the translated address falls within a reserved area of the medium (e.g. speed tolerance gap, alternate logical block, vendor reserved area, etc.). If the entire translated address falls within a reserved area the target may not return a translated address. An RAREA bit of zero indicates that no part of the translated address falls within a reserved area of the medium.

Table 5.1.1-14a. Address Field Logical Block Address Format

Bit Byte(s)	7	6	5	4	3	2	1	0
0	Logical Block Address							
1								
2								
3								
4	0	0	0	0	0	0	0	0
5	0	0	0	0	0	0	0	0
6	0	0	0	0	0	0	0	0
7	0	0	0	0	0	0	0	0

Table 5.1.1-14b. Address Field Physical Sector Address Format

Bit Byte(s)	7	6	5	4	3	2	1	0
0	Cylinder Number							
1								
2								
3	Head Number							
4	Sector Number							
5								
6								
7								

- [5] An alternate sector (ALTSEC) bit of one indicates that the translated address is physically located in an alternate sector of the medium. If the drive cannot determine if all or part of the translated address is located in an alternate sector it shall set this bit to zero. An ALTSEC bit of zero indicates 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.
- [6] An alternate track (ALTTRK) bit of one indicates 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. An ALTTRK bit of zero indicates that no part of the translated address is located on an alternate track of the medium.
- [7] The Translated Format field contains the value from the Send Diagnostic command translate format field (see Table 5.1.1-17). The values are 000 (Logical block format) or 101 (Physical sector address format).
- [8] The Translated Address field contains the address the target translated from the address supplied by the initiator in the Send Diagnostic command. This field shall be in the format specified in the translate format field. The supported formats are shown in Tables 5.1.1-14a and 5.1.1-14b.

5.1.1.6 Send Diagnostic Command (1Dh)

Table 5.1.1-15. Send Diagnostic Command

Bit Byte(s)	7	6	5	4	3	2	1	0
0	0	0	0	1	1	1	0	1
1	Logical Unit No. 0 0 0			PF[1] 0	RSVD	SELF TEST [2]	DEVOFL [3]	UnitOFL [4]
2	0	0	0	0	0	0	0	0
3	Parameter List Length (MSB) 0 0 0 0 0 0 0 0							
4	Parameter List Length (LSB) 0 0 0 0 (see explanation below) [5]							
5	0	0	0	0	0	0	Flag	Link

This command requests that the drive perform diagnostic tests on itself, or perform other optional operations. Table 5.1.1-15 shows the format of the Send Diagnostics command as implemented by the drive. When the Self Test bit is zero (see note [2]), this command is usually followed by a Receive Diagnostics Results command and a subsequent data-in phase that returns data to the initiator. Using this latter procedure of Send Diagnostics/Receive Diagnostics Results commands the initiator can ask the drive to return a list of optional operations it supports (Table 5.1.1-16) and then request a supported additional operation. The drive supports only the Translate Address operation (Table 5.1.1-17).

Translate Address Page - Send Diagnostic

The translate address page allows the initiator to translate a logical block address into a physical sector address or a physical sector into a logical block address. The address to be translated is passed to the drive with Send Diagnostic command and the results are returned to the initiator during the data in phase following the Receive Diagnostic Results command. The format of the translate address page - Send Diagnostic is shown in Table 5.1.1-17. The translated address is returned in the translate address page returned after the Receive Diagnostic Results Command (see Table 5.1.1-11).

Table 5.1.1-17. Translate Address Page - Send Diagnostic

Bit Byte(s)	7	6	5	4	3	2	1	0	
0	Page Code (40h)								
1	Reserved								
2	(MSB)	Page Length (000Ah)							
3							(LSB)		
4	Reserved					Supplied Format [1]			
5	Reserved					Translate Format [2]			
6 : 13	Address To Translate [3]								

Notes.

- [1] The Supplied Format field specifies the format of the address to translate field. The valid values for this field are 000 for logical block address format or 101 for physical sector address format. If the drive does not support the requested format it shall terminate the Send Diagnostic command with Check Condition status. The sense key shall be set to Illegal Request and an additional sense code shall be set to Invalid Field In Parameter List.
- [2] The Translate Format field specifies the format to which the initiator would like the address to be translated. The valid values for this field are 000 for logical block address format or 101 for physical sector address format. The Translate Format field must be different than the Supplied Format Field. If the drive does not support the requested format it shall terminate the command with Check Condition status. The sense key shall be set to Illegal Request and an additional sense code shall be set to Invalid Field In Parameter List.
- [3] The Address to Translate field contains a single address the initiator is requesting the drive to translate. The format of this field is defined by the Supplied Format Field. The supported formats are shown in Table 5.1.1-14a and Table 5.1.1-14b.

For systems which support disconnection, the drive will disconnect while executing this command.

5.1.2 Group 1 Commands for all device types

The following commands are applicable for all device types, but the drive only implements those indicated in the paragraphs following.

OP Code	Command Name	Section	Page
39h	Compare Command	5.1.2.1	108
3Ah	Copy and Verify Command	5.1.2.2	108
3Bh	Write Buffer	5.1.2.3	108
3Ch	Read Buffer	5.1.2.4	113

5.1.2.1 Compare command (39h)

Not Implemented. A "Check Condition" Status is sent if received.

5.1.2.2 Copy and Verify command (3Ah)

Not Implemented. A "Check Condition" status is sent if received.

5.1.2.3 Write Buffer command (3Bh)

Table 5.1.2-1. Write Buffer Command

Bit Byte(s)	7	6	5	4	3	2	1	0
0	0	0	1	1	1	0	1	1
1	Logical Unit No. [1]			RSVD		Mode [2]		
2	Buffer ID							[5]
3	(MSB)							[6]
4	Buffer Offset							[6]
5								[6] (LSB)
6	(MSB)							
7	Byte Transfer Length [3]							
8								(LSB)
9	0	0	0	0	0	0	Flag [4]	Link [4]

The Write Buffer command may be used in conjunction with the Read Buffer command as a diagnostic function for testing the drive's data buffer memory and the SCSI bus integrity. When used in a diagnostic mode, the medium shall not be accessed during the execution of this command. Additional modes are provided for downloading and saving executable micro-code.

The function of this command and the meaning of the fields within the Command Descriptor Block depend on the mode field (Byte 1, data bits 0, 1, 2). See note [2].

Notes.

- [1] LUN must be zero
 [2] The mode field is defined in the following table and in the referenced sections.

Check with your drive Product Manual Volume 1 to see which modes are supported by the drive in question.

DB2	DB1	DB0	Mode Definition
0	0	0	Write combined header and data (section 5.1.2.3.1)
0	1	0	Write data (section 5.1.2.3.2)
1	0	0	Download microcode (section 5.1.2.3.3)
1	0	1	Download microcode and save (section 5.1.2.3.4)
1	1	0	Download microcode with offsets
1	1	1	Download microcode with offsets and save

- [3] Use a Read Data Buffer command with mode bits set to 011b to get the drive buffer capacity to use with the Write Buffer command.
 [4] See "Control Byte" paragraph 4.2.6. All bits are zero for all modes.
 [5][6] Implemented only by drives that support modes 110b and 111b.

5.1.2.3.1 Combined Header and Data Mode (000b)

In this mode, data to be written to the drives data buffer is preceded by a four byte header.

The Byte Transfer Length includes a four byte header and the Write Buffer data. A transfer length of zero indicates that no data transfer shall take place. This condition shall not create the Check Condition status. If the transfer length is greater than the Buffer Capacity reported by the Read Buffer header, the drive shall create the Check Condition status with the Sense Key of Illegal Request. In this case no data shall be transferred from the initiator.

It shall not be considered an error to request a transfer length less than the Buffer Capacity.

Buffer ID and Buffer offset fields are all zero.

The write data following the Write Buffer CDB consists of a 4 byte write buffer header (which always precedes the data) plus the data to be written to the data buffer as follows:

Table 5.1.2-2. Write Buffer Header

Bit Byte(s)	7	6	5	4	3	2	1	0
0	0	0	0	0	0	0	0	0
1	0	0	0	0	0	0	0	0
2	0	0	0	0	0	0	0	0
3	0	0	0	0	0	0	0	0
4 - n	Data To Be Written Into Drive Buffer							

5.1.2.3.2 Write data only mode (010b)

The byte transfer length specifies the maximum number of bytes that shall be transferred during the Data Out phase to be stored in the drive buffer. No header bytes are included. The buffer ID and buffer offset fields are all zero.

5.1.2.3.3 Download Microcode mode (100b)

If the logical unit cannot accept this command because of some device condition, the logical unit shall terminate each WRITE BUFFER command with this mode (100b) with a CHECK CONDITION status, a sense key of ILLEGAL REQUEST, and shall set the additional sense code to COMMAND SEQUENCE ERROR.

In this mode, vendor-specific microcode or control information shall be transferred to the control memory space of the logical unit. After a power-cycle or reset, the device operation shall revert to a vendor-specific condition. The meanings of the buffer ID, buffer offset, and parameter list length fields are not specified by the International Standard and are not required to be zero-filled. When the microcode download has completed successfully the logical unit shall generate a unit attention condition for all initiators except the one that issued the WRITE BUFFER command. The additional sense code shall be set to MICROCODE HAS BEEN CHANGED.

5.1.2.3.4 Download and Save Microcode mode (101b)

In this mode, vendor-unique executable microcode (which is not preceded by a 4 byte header) shall be transferred to the control memory space of the target and, if the download is completed successfully, shall also be saved. The downloaded code shall then be effective after each power cycle and reset until it is supplanted in another download microcode and save operation. When the download microcode and save command has been completed successfully the target shall generate a "Unit Attention Condition" (see section 4.6) for all initiators with an extended sense of "Power On, Reset, or Bus Device Reset Occurred" (Sense data error code 29 00h). Following the downloading of new microcode the drive may need to be reformatted before it can perform properly.

For this mode (101b) the command bytes of Table 5.1.2-1 are interpreted as shown below:

Buffer ID and Buffer offset fields (CDB bytes 2-5):

These bytes are all zero.

Byte Transfer Length (CDB bytes 6, 7 and 8):

The transfer length in bytes of the downloadable code. This value must be the exact length of the download data. A value of one signifies one byte of download data, etc.

5.1.2.3.5 Download microcode with offsets (110b)

In this mode, the initiator may split the transfer of the vendor-specific microcode or control information over two or more WRITE BUFFER commands. If the logical unit cannot accept this command because of some device condition, the logical unit shall terminate each WRITE BUFFER command with this mode (110b) with a CHECK CONDITION status, a sense key of ILLEGAL REQUEST, and shall set the additional sense code to COMMAND SEQUENCE ERROR.

If the last WRITE BUFFER command of a set of one or more commands completes successfully, the microcode or control information shall be transferred to the control memory space of the logical unit. After a power-cycle or reset, the device shall revert to a vendor-specific condition. In this mode, the DATA OUT phase contains vendor-specific, self-describing microcode or control information.

Since the downloaded microcode or control information may be sent using several commands, when the logical unit detects the last download microcode with offsets and save mode WRITE BUFFER command has been received, the logical unit shall perform any logical unit required verification of the complete set of downloaded microcode or control information prior to returning GOOD status for the last command. After the last command completes successfully the logical unit shall generate a unit attention condition (see Section 4.6) for all initiators except the one that issued the set of WRITE BUFFER commands. When reporting the unit attention condition, the logical unit shall set the additional sense code to MICROCODE HAS BEEN CHANGED.

If the complete set of WRITE BUFFER commands required to effect a microcode or control information change (one or more commands) are not received before a reset or power-on cycle occurs, the change shall not be effective and the microcode or control information is discarded.

The buffer ID field identifies a specific buffer within the logical unit. The vendor assigns buffer ID codes to buffers within the logical unit. A Buffer ID value of zero shall be supported. If more than one buffer is supported, additional buffer ID codes shall be assigned contiguously, beginning with one. If an unsupported buffer ID code is identified, the logical unit shall return CHECK CONDITION status and shall set the sense key to ILLEGAL REQUEST with an additional sense code of INVALID FIELD IN CDB.

The microcode or control information are written to the logical unit buffer starting at the location specified by the buffer offset. The initiator shall conform to the offset boundary requirements. If the logical unit is unable to accept the specified buffer offset, it shall return CHECK CONDITION status and it shall set the sense key to ILLEGAL REQUEST with an additional sense code of INVALID FIELD IN CDB.

The parameter list length specifies the maximum number of bytes that shall be transferred during the DATA OUT phase to be stored in the specified buffer beginning at the buffer offset. The initiator should attempt to ensure that the parameter list length plus the buffer offset does not exceed the capacity of the specified buffer. (The capacity of the buffer can be determined by the buffer capacity field in the READ BUFFER descriptor.) If the buffer offset and parameter list length fields specify a transfer that would exceed the buffer capacity, the logical unit shall return CHECK CONDITION status and shall set the sense key to ILLEGAL REQUEST with an additional sense code of INVALID FIELD IN CDB.

5.1.2.3.6 Download microcode with offsets and save mode (111b)

In this mode, the initiator may split the transfer of the vendor-specific microcode or control information over two or more WRITE BUFFER commands. If the logical unit cannot accept this command because of some device condition, the logical unit shall terminate each WRITE BUFFER command with this mode (111b) with a CHECK CONDITION status, a sense key of ILLEGAL REQUEST, and shall set the additional sense code to COMMAND SEQUENCE ERROR.

If the last WRITE BUFFER command of a set of one or more commands completes successfully, the microcode or control information shall be saved in a nonvolatile memory space (semiconductor, disk or other). The saved downloaded microcode or control information shall then be effective after each power-cycle and reset until it is supplanted by another download microcode with save operation or download microcode with offsets and save operation. In this mode, the DATA OUT phase contains vendor-specific, self-describing microcode or control information.

Since the downloaded microcode or control information may be sent using several commands, when the logical unit detects the last download microcode with offsets and save WRITE BUFFER command has been received, the logical unit shall perform any logical unit required verification of the complete set of downloaded microcode or control information prior to returning GOOD status for the last command. After the last command completes successfully the logical unit shall generate a unit attention condition (see Section 4.6) for all initiators except the one that issued the set of WRITE BUFFER commands. When reporting the unit attention condition, the logical unit shall set the additional sense code to MICROCODE HAS BEEN CHANGED.

If the complete set of WRITE BUFFER commands required to effect a microcode or control information change (one or more commands) are not received before a reset or power-on cycle occurs, the change shall not be effective and the microcode or control information is discarded.

The buffer ID field identifies a specific buffer within the logical unit. The vendor assigns buffer ID codes to buffers within the logical unit. A Buffer ID value of zero shall be supported. If more than one buffer is supported, additional buffer ID codes shall be assigned contiguously, beginning with one. If an unsupported buffer ID code is identified, the logical unit shall return CHECK CONDITION status and shall set the sense key to ILLEGAL REQUEST with an additional sense code of INVALID FIELD IN CDB.

The microcode or control information are written to the logical unit buffer starting at the location specified by the buffer offset. The initiator shall conform to the offset boundary requirements. If the logical unit is unable to accept the specified buffer offset, it shall return CHECK CONDITION status and it shall set the sense key to ILLEGAL REQUEST with an additional sense code of INVALID FIELD IN CDB.

The parameter list length specifies the maximum number of bytes that shall be transferred during the DATA OUT phase to be stored in the specified buffer beginning at the buffer offset. The initiator should attempt to ensure that the parameter list length plus the buffer offset does not exceed the capacity of the specified buffer. (The capacity of the buffer offset. The initiator should attempt to ensure that the parameter list length plus the buffer offset does not exceed the capacity field in READ BUFFER descriptor.) If the buffer offset and parameter list length fields specify a transfer that would exceed the buffer capacity, the logical unit shall return CHECK CONDITION status and shall set the sense key to ILLEGAL REQUEST with an additional sense code of INVALID FIELD IN CDB.

5.1.2.4 Read Buffer command (3Ch)

Table 5.1.2-3. Read Buffer

Bit Byte(s)	7	6	5	4	3	2	1	0
0	0	0	1	1	1	1	0	0
1	Logical Unit No [1] 0 0		0	RSVD 0 0		Mode [2] 0 0		0
2	Buffer ID [5]							
3	(MSB) _____							[6]
4	Buffer Offset							[6]
5								[6](LSB)
6	(MSB) _____							(LSB)
7	Allocation Length [3]							
8								
9	0	0	0	0	0	0	Flag [4]	Link [4]

The Read Buffer command is used in conjunction with the Write Buffer command as a diagnostic function for testing drive memory and the SCSI bus integrity. This command shall not alter the medium. Command format is shown in Table 5.1.2-3.

Notes. (Notes [3], [4], [5] and [6] appear on page 115 after paragraph 5.2.1.4.3)

[1] LUN must be zero.

[2] Read Buffer Mode Field. The interpretation of data bits 0-2 is given in the following table and in paragraphs following the table.

DB2	DB1	DB0	
0	0	0	Read combined descriptor header and data (section 5.1.2.4.1)
0	1	0	Read data (section 5.1.2.4.2)
0	1	1	Read descriptor (section 5.1.2.4.3)

**5.1.2.4.1 Read Combined Descriptor Header and Data mode (000b)
(see Table 5.1.2-4)**

In this mode, a four byte Read Buffer header followed by the data bytes are returned to the initiator during the Data In phase. The buffer ID and the buffer offset fields are not used by drives supported by this manual, and must be zero.

Table 5.1.2-4. Read Buffer Header (Mode 000b)

Bit Byte(s)	7	6	5	4	3	2	1	0
0	0	0	0	0	0	0	0	0
1	(MSB) _____							(LSB)
2	Buffer Capacity							
3								

continued from previous page

5.1.2.4.2 Read data

In this mode, the DATA IN phase contains buffer data only with no header. The buffer ID and buffer offset fields are not used.

5.1.2.4.3 Read Buffer descriptor mode (011b)

In this mode, a maximum of four bytes of READ BUFFER descriptor information are returned. If there is no buffer associated with the specified buffer ID, the target shall return all zeros in the READ BUFFER descriptor. The buffer offset field is reserved in this mode. The allocation length should be set to four or greater. The target shall transfer the lesser of the allocation length or four bytes of READ BUFFER descriptor. The READ BUFFER descriptor is defined as shown in the Table following.

Table 5.1.2.4.3-1. READ BUFFER descriptor

Bit Byte(s)	7	6	5	4	3	2	1	0	
0	Offset Boundary								
1	(MSB)	Buffer Capacity							
2									
3								(LSB)	

The offset boundary field returns the boundary alignment with the selected buffer for subsequent WRITE BUFFER and READ BUFFER commands. The value contained in the offset boundary field shall be interpreted as a power of two.

The value contained in the buffer field of subsequent WRITE BUFFER and READ BUFFER commands should be a multiple of $2^{\text{offset boundary}}$ as shown in table following.

Table 5.1.2.4.3-2 Buffer Offset boundary

Offset Boundary	$2^{\text{Offset Boundary}}$	Buffer Offsets
0	$2^0 = 1$	Byte boundaries
1	$2^1 = 2$	Even-byte boundaries
2	$2^2 = 4$	Four-byte boundaries
3	$2^3 = 8$	Eight-byte boundaries
4	$2^4 = 16$	16-byte boundaries
.	.	.
.	.	.
FFh	Not applicable	0 is the only supported buffer offset.

The buffer capacity field shall return the size of the selected buffer in bytes.

IMPLEMENTORS NOTE: In a multi-tasking system, a buffer may be altered between the WRITE BUFFER and READ BUFFER commands by another task. Buffer testing applications may wish to insure that only a single task is active. Use of reservations (to all logical units on the device) may also be helpful in avoiding buffer alteration between these two commands.

continued from previous page

- [3] Drive specific size also depends on whether cache is enabled as shown in Table 5.2.1-27 (RCD = 0) or not (RCD = 1). See drive Product Manual under list of SCSI commands supported.

The buffer capacity field specifies the total number of data bytes that are available in the drive's data buffer (see 5.1.2.3.1 and 5.1.2.3.2). This number is not reduced to reflect the allocation length nor is it reduced to reflect the actual number of bytes written using the Write Buffer command. Following the Read Buffer header, the drive shall transfer data from its data buffer. The drive terminates the Data In phase when allocation length bytes of header plus data have been transferred or when all available header and buffer data have been transferred to the initiator, whichever is less.

- [4] See "Control Byte" paragraph 4.2.6.

- [5][6] Not implemented by drives supported by this manual. Must be zero.

5.1.3 Group 2 commands for all device types

This group consists of the following supported 10 byte commands:

5.1.3.1 Change Definition Command (40h)

5.1.3.2 Log Select Command (4Ch)

5.1.3.3 Log Sense Command (4Dh)

5.1.3.1 Change Definition command (40h)

Table 5.1.3-1. Change Definition command

Bit Byte(s)	7	6	5	4	3	2	1	0
0	0	1	0	0	0	0	0	0
1	Logical Unit No. [1]			Reserved				
2	Reserved							Save [2]
3	RSVD	Definition Parameter [3]						
4	Reserved							
5	Reserved							
6	Reserved							
7	Reserved							
8	Parameter Data Length [4]							
9	Control [5]							

The Change Definition command (Table 5.1.3-1) modifies the operating definition of the drive with respect to commands from all initiators.

The drive maintains only one operating definition and it applies to *all* initiators in the system.

Notes.

- [1] The LUN must be zero.

- [2] A save control bit (Save) of zero indicates that the drive shall not save the operating definition. A Save bit of one indicates that the drive shall save the operating definition to non-volatile memory.

- [3] The definition parameter field is defined in Table 5.1.3-2.

Table 5.1.3-2. Definition Parameter Field

Value	Meaning of Definition Parameter
00h	Use Current Operating Definition
01h	SCSI X3.131-1986 Operating Definition
02h*	CCS Operating Definition
03h	SCSI-2 X3.131-198X Operating Definition
04-3Fh	Reserved
40-7Fh	Vendor Specific

*The drive treats an 02 option as if it had been 01.

The current operating definition parameter values establish operating definitions compatible with the appropriate SCSI specification.

[4] The parameter data length field specifies the length in bytes of the parameter data that shall be transferred from the initiator to the target. A parameter data length of zero indicates that no data shall be transferred. This condition shall not be considered as an error. Parameter data lengths greater than zero indicate the number of bytes of parameter data that shall be transferred.

The parameter data is not used by the drive.

[5] Control is not used by the drive.

The operating definition is modified after successful completion of the command. The drive shall consider the command successfully completed when it detects the assertion of the ACK signal for the Command Complete message. The initiator should verify the new operating definition by issuing an Inquiry command requesting the implemented operating definition page (see Table 5.1.1-10d).

It is permissible for a SCSI-2 device that has its definition changed to a SCSI-1 device to accept a Change Definition command.

If the Change Definition command is not executed successfully for any reason, the operating definition shall remain the same as it was before the Change Definition command was attempted.

After a power-on condition or a hard Reset condition, the drive shall set its initial operating definition to the last saved value, if saving is implemented, or its default value, if saving is not implemented. Default is SCSI 2 mode.

5.1.3.2 Log Select command (4Ch)

The Log Select command provides a means for an initiator to manage statistical information about the drive operation. This information is logged within the drive and can be sent to the initiator in response to a Log Sense command from the initiator. The Log Select command format is shown in Table 5.1.3.2-1. In the Data Out phase following the command, the initiator sends zero or more pages of control parameters in the Log Page Format of Table 5.1.3.2-3. These Log Pages contain parameters that command the drive to change selected threshold, or cumulative values of any or all drive logs. Numbers in brackets [] refer to notes following tables.

The following tables in section 5.1.3.2 apply for the Log Select command as indicators of functions that command the drive to perform or enable for performance, control parameter bits the drive shall set/reset/save, log counts that shall be kept, etc. For the Log Sense command [1] these tables apply as indicators of functions the drive reports back to the host that it is enabled to perform, control parameter bits that are set/reset/saved, log counts that are being kept, etc. Though the language of the descriptions is for the Log Select case, the application to the Log Sense case should also be considered.

The drives represented by this Interface Manual do not support keeping independent sets of log parameters (one set for each initiator in the system). If at some point log parameters are changed (by a Log Select command) that affect initiators other than the initiator that sent the Log Select command, the drive generates a unit attention condition for those other initiators, but not for the one that issued the Log Select command. When the other initiators at a future time connect to the drive, the first command attempted would not execute and a check condition status would be issued by the drive. A Request Sense command would normally follow and a unit attention condition sense code be returned to these other initiators with an additional sense code of LOG PARAMETERS CHANGED. (one by one as they connect to the drive) See Section 4.6, Unit Attention Condition.

[1] Section 5.1.3.3 describes the Log Sense command, but the tables of this section that apply are not repeated there.

Table 5.1.3.2-1. The Log Select command (4Ch)

Bit Byte(s)	7	6	5	4	3	2	1	0	
0	0	1	0	0	1	1	0	0	
1	Logical Unit Num 0 0 0			Reserved			PCR [1]	SP [2]	
2	PC [3]		Reserved						
3	Reserved								
4	Reserved								
5	Reserved								
6	Reserved								
7	(MSB)				Parameter List Length [4]				
8									(LSB)
9	0	0	0	0	0	0	Flag [5]	Link [5]	

[1] **Parameter Code Reset (PCR).** A PCR bit of one and a Parameter List Length of zero causes all implemented parameters to be set to the drive-defined default values (most likely zero). If the PCR bit is one and the parameter list length is greater than zero, the command is terminated with a Check Condition Status. The sense key shall be set to Illegal Request and the additional sense code shall be set to Invalid Field in CDB. A PCR bit of Zero specifies that the log parameters shall not be reset.

[2] **Save Parameters (SP) bit.** An SP bit of one indicates that after performing the specified Log Select operation the drive shall save to non-volatile memory all Log Select parameters identified as savable by the DS bit in the Log parameter sections of the Log Page (see Tables 5.1.3.2-3 and 5.1.3.2-4). An SP bit of zero specifies that parameters shall not be saved. Log Parameters are also saved after each thermal calibration if the TSD bit in the Log Parameter pages (see Table 5.1.3.2-4) is zero.

It is not an error to set the SP bit to one and to set the DS bit of a log parameter to one. In this case, the parameter value for that log parameter is not saved.

[3] **Page Control field (PC).** This field defines the type of log parameter the initiator selects to change with the Log Select Command. The PC field is defined in Table 5.1.3.2-2.

Table 5.1.3.2-2. Page Control Field

PC Field Value	Type of Log Parameter
00b	Log Current Threshold Values
01b	Log Current Cumulative Values
10b	Log Default Threshold Values
11b	Log Default Cumulative Values

The drive only updates the cumulative values to reflect the number of events experienced by the drive, but the initiator can set the threshold or cumulative log (00 or 01) parameter values using the Log Select command with the PC field set as applicable.

The drive sets the current log values to default values in response to a Log Select command with the parameter list length set to zero and the PC field set to the applicable value (10 or 11) per Table 5.1.3.2-2.

continued from previous page

If an initiator attempts to change a current threshold value that is not available or not implemented for that log parameter, the drive shall terminate the Log Select command with a Check Condition status, the sense key set to Illegal Request and an additional sense code set to Invalid Field In Parameter List. The saving of current threshold parameters and the criteria for the current threshold being met are controlled by bits in the Parameter Control Byte (PCB) (byte 2 of each of the Log Parameter pages).

- [4] Parameter List Length. This specifies the length in bytes of the parameter list that shall be transferred from the initiator to the drive during the Data Out phase. A parameter list length of zero indicates that no pages shall be transferred. This condition shall not be considered an error.

If the initiator sends page codes or parameter codes within the parameter list that are reserved or not implemented by the drive, the drive shall terminate the Log Select command with Check Condition status. The sense key shall be set to Illegal Request and the additional sense code set to Invalid Field In Parameter List.

If a parameter list length results in the truncation of any log parameter, the drive shall terminate the command with Check Condition status. The sense key shall be set to Illegal Request and the additional sense code set to Invalid Field In CDB.

In the Data Out phase of the Log Select command the initiator may send none, one, or more data pages, each of which is in the format specified in Table 5.1.3.2-3 and which contain control information pertaining to the management and reporting of various drive log parameters. If multiple pages are sent out following the command CDB, they must be sent in ascending page code value order. Also, Log Parameters in each Log Page must be sent in Log Parameter Code ascending order (see Table 5.1.3.2-3 Note 4). The Page Code (byte 0) specifies which log the page pertains to. The Page Code assignments are listed in Table 5.1.3.2-6. Initiator cannot send page zero.

- [5] Flag and Link bits used as defined in paragraph 4.2.6.

The Log Sense command and its data-in phase also use the format of Table 5.1.3.2-3 so explanations pertaining to the table apply generally to both Log Select and Log Sense commands, with differences noted. However, only one log page is selected and returned with each Log Sense command. See Section 5.1.3.3. This Table is not repeated in Section 5.1.3.3.

Table 5.1.3.2-3. Log Page Format

Bit Byte(s)	7	6	5	4	3	2	1	0
0	Reserved		Page Code [1]					
1	Reserved							
2	(MSB)		Page Length (n-3) [2]				(LSB)	
3								

Log Parameter Structure(s) [3]

4 to x+3	Log Parameter (First) [4] (Length X bytes)							
.	.							
.	.							
.	.							
n-Y to	Log Parameter (last) [4] (Length Y bytes)							

Notes for Table 5.1.3.2-3.

- [1] Page codes implemented are given in Table 5.1.3.2-6.
- [2] Gives the total number of bytes of Log Parameter structures that follow these first four control block bytes. If the initiator sends a page length that results in the truncation of any parameter, the target shall terminate the command with Check Condition status. The sense key shall be set to Illegal Request with the additional sense code set to Invalid Field In Parameter List.
- [3] Most log pages contain one or more special data structures called Log Parameters. Log Parameters may be data counters that record a count of a particular event (or events) or list parameters (strings) that contain a description of a particular event. List parameters are not currently supported by the drives represented by this manual.
- [4] Each Log Parameter structure begins with a four byte parameter header followed by one or more bytes of parameter value data. Log Parameter structures are in the format given in Table 5.1.3.2-4.

Table 5.1.3.2-4. Log Parameters

Bit Byte(s)	7	6	5	4	3	2	1	0
0	(MSB) Parameter Code [1]							
1	(LSB)							
2	DU [2]	DS [3]	TSD [4]	ETC [5]	TMC [6]		RES- RVD	LP [7]
3	Parameter Length (n-3 bytes)							[8]
4	--	Parameter Value [9]						--
:	--							--
n								

Notes for Table 5.1.3.2-4

- [1] The parameter code field identifies the specific parameter that is being transferred with the Log Page. These codes are listed and explained in the individual page code descriptions following Table 5.1.3.2-6.

Byte 2 is referred to as the parameter control byte. The meanings of the various fields are discussed in notes [2] through [7]. For a Log Select command these bits perform a control function, but on a Log Sense command they only report the drive settings of these bits in this same format on the data-in part of the Log Sense command.

- [2] Disable Update (DU). For the Log Select command, this applies only to the cumulative log parameter values (indicated by 01 in the PC field of the Log Select and Log Sense command descriptor block). The drive can set this bit to one or zero also (see note [7] following).

The DU flag bit is defined as follows:

- (a) A zero value indicates that the drive shall update the log parameter value to reflect all events that should be logged by that parameter.
- (b) A one value indicates that the drive shall not update the log parameter value except in response to a Log Select command that specifies a new value for the parameter.

The DU flag is set to one when the current cumulative value of the parameter counter it controls reaches its maximum value (see Note 8). Upon reaching this maximum value, the data counter does not wrap around and start over at zero. Incrementing of other counters within the same log pages ceases. Counters do not restart automatically if the overflowed counter is re-initialized. If the data counter reaches its maximum value during the execution of a command, the drive completes the command. Drive counter updates are performed in the background. This means a counter may overflow long after a command has completed, so the drive must treat this condition as a Unit Attention with the Additional Sense code set to Log Counter at max for all initiators if RLEC=1 (Report Log exception condition bit of the Control Mode Page 0Ah).

Since the drive uses volatile memory to hold cumulative values, they will be lost when a power cycle occurs. Unless the initiator commands the drive to save them to non-volatile memory using a Log Select or Log Sense command with the SP bit set to one.

continued from previous page

The DU bit is not defined for threshold values (indicated by the PC field of the Log Sense command descriptor block) nor for list parameters (indicated by the LP bit). The drive ignores the value of DU bits in a Log Select command applicable to threshold values or list parameters.

- [3] Disable Save (DS). If DS is zero it indicates that the drive shall support (Log Select)/supports (Log Sense data) saving for that log parameter. The drive shall save/does save the current cumulative and the current threshold parameter values in response to a Log Select or Log Sense command with a SP bit of one. A DS bit of one indicates that the drive shall not/does not support saving that log parameter in response to a Log Select or Log Sense command with a SP bit of one.
- [4] Target Save Disable (TSD). A zero indicates that the drive shall save frequently enough to insure statistical significance. The drive's method is to save after each thermal calibration, which is once every ten minutes. A one bit indicates that the drive shall not/does not use its save method.
- [5] Enable Threshold Comparison (ETC). A one indicates that a comparison to the threshold value shall be (Log Select)/is (Log Sense data) performed whenever the cumulative value is updated. A bit of zero indicates the comparison shall not be/is not performed. The value of the ETC bit is the same for both the threshold and cumulative parameters.
- [6] Threshold Met Criteria (TMC). This field defines the basis for comparison of the cumulative and threshold values. See Table 5.1.3.2-5 for meanings of values in this field. The TMC field is only valid when the ETC bit is one.

Table 5.1.3.2-5. Threshold Met Criteria

Code	Basis for Comparison
00b	Notify of every update of cumulative value
01b*	Cumulative value equal to threshold value
10b*	Cumulative value not equal threshold value
11b*	Cumulative value greater than threshold value

*Comparison made at every update of cumulative value.

If the ETC bit is one* and the result of the comparison is true, a unit attention condition is generated for all initiators. When reporting the unit attention condition the drive sets the sense key to Unit Attention, and the additional sense code to Threshold Condition Met.

*The RLEC bit (Report Log Exception Condition) in Mode page 0AH (Table 5.2.1-28) must also be one.

- [7] List Parameter (LP). Zero indicates the parameter is a data counter. One indicates that the parameter is a list parameter. This bit only has meaning for the Log Sense command data-in pages.

Data counters are associated with one or more events. The data counter is updated whenever one of these events occurs by incrementing the counter value, provided the DU bit is zero. See Note 2 above.

An LP bit of one indicates that the parameter is a list parameter. List parameters are not counters and thus the ETC and TMC fields shall be set to zero. A list parameter is a string of ASCII graphic codes (i.e., code values 20h thru 73h). List parameters are not supported by the drive at this time.

continued from previous page

- [8] **Parameter Length.** This field specifies the length in bytes of the parameter that follows. If the initiator sends a parameter length value that results in the truncation of the parameter value, the drive terminates the command with a Check Condition status. The sense key shall be set to Illegal Request with the additional sense code set to Invalid Field In Parameter List.
- [9] **Parameter Value.** This field uses one, two, four or eight bytes to transmit an unsigned counter value. The initiator sends counts to set values into counters in the drive, and the drive returns counter values to the initiator. The initiator is responsible to issue a Log Sense command to learn the parameter length the target has selected.

When any counter in a log page reaches its maximum value, the drive ceases incrementing all counters in that log page. If the RLEC bit of the Control Mode page is one, then the drive reports the exception condition as described in note [2] above.

The page code assignments for the log pages are listed in Table 5.1.3.2-6. Detailed descriptions follow the table.

Table 5.1.3.2-6. Log Page Codes

Page Code	Description	Section	Page
02h	Error Counter Page (Write)	5.1.3.2.1	123
03h	Error Counter Page (Read)	5.1.3.2.1	123
05h	Error Counter Page (Verify)	5.1.3.2.1	123
06h	Non-Medium Error page	5.1.3.2.2	125
37h	Cache Statistics page	5.1.3.2.3	125
3Eh	Factory Log page	5.1.3.2.4	126

5.1.3.2.1 Write, Read and Verify Error Counter pages (code 02, 03, 05h)

Table 5.1.3.2-7 defines the parameter code field for the write, read and verify error counter pages.

Table 5.1.3.2-7. Write, Read and Verify Errors Parameter Code Field

Bit Byte(s)	7	6	5	4	3	2	1	0
0	0	0	0	0	0	0	0	0
1	Parameter Code [1]							

Notes.

- [1] Parameter codes 00h through 06h specify six counters each for write, read and verify errors (18 counters). A description of the type (category of error) counters specified by codes 00h through 06h are described following.

Parameter Code 00h - Error Corrected Without Substantial Delay. An error correction was applied to get perfect data (a.k.a. ECC on-the fly). 'Without Substantial Delay' means the correction did not postpone reading of later sectors (e.g., a revolution was not lost). The counter is incremented once for each logical block that requires correction. Two different blocks corrected during the same command are counted as two events.

Parameter Code 01h - Error Corrected with Possible Delays. An error code or algorithm (e.g., ECC, checksum) is applied in order to get perfect data with substantial delay. "With possible delay" means the correction took longer than a sector time so that reading/writing of subsequent sectors was delayed (e.g., a lost revolution). The counter is incremented once for each logical block that requires correction. A block with a double error that is correctable counts as one event and two different blocks corrected during the same command count as two events.

Parameter Code 02h - Total (e.g., re-writes or re-reads) This parameter code specifies the counter counting the number of errors that are corrected by applying retries. This counts errors recovered, not the number of retries. If five retries were required to recover one block of data, the counter increments by one, not five. The counter is incremented once for each logical block that is recovered using retries. If an error is not recoverable while applying retries and is recovered by ECC, it isn't counted by this counter; it will be counted by the counter specified by parameter code 01h - Error Corrected with Possible Delay.

Parameter code 03h - Total Error Corrected. This counter counts the total of parameter code errors 00h, 01h and 02h. There is to be no "double counting" of data errors among these three counters. The sum of all correctable errors can be reached by adding parameter code 01h and 02h errors, not by using this total.

Parameter Code 04h - Total Times Correction Algorithm Processed. This parameter code specifies the counter that counts the total number of retries, or "times the retry algorithm, is invoked". If after five attempts a counter 02h type error is recovered, then five is added to this counter. If three retries are required to get a stable ECC syndrome before a counter 01h type error is corrected, then those three retries are also counted here. The number of retries applied to unsuccessfully recover an error (counter 06h type error) are also counted by this counter.

Parameter Code 05h - Total Bytes Processed. This parameter code specifies the counter that counts the total number of bytes either successfully or unsuccessfully read, written or verified (depending on the log page) from the drive. If a transfer terminates early because of an unrecoverable error, only the logical blocks up to and including the one with the unrecoverable error are counted. Data bytes transferred to the initiator during a Mode Select, Mode Sense, Inquiry, Write Data Buffer, etc. do not count; only user data bytes are counted by this counter.

Parameter Code 06h - Total Uncorrected Errors. This parameter code specifies the counter that contains the total number of blocks for which an uncorrected data error has occurred.

5.1.3.2.2 Non-Medium Error Page (code 06h)

Log page code 06h specifies non-medium errors. The page format is shown in Table 5.1.3.2-8.

Table 5.1.3.2-8. Non-Medium Error Page (Code 06h)

Bit Byte(s)	7	6	5	4	3	2	1	0
0	0	0	0	0	0	0	0	0
1	Parameter Code [1]							

Notes.

[1] Parameter code 00h is the only code supported for this page and it represents the number of recoverable error events other than write, read or verify errors.

5.1.3.2.3 Cache Statistics Page (Page code 37h)

Log Page code 37h specifies Cache Statistics page. The page format is shown in Table 5.1.3.2-9.

Table 5.1.3.2-9. Cache Statistics Page

Bit Byte(s)	7	6	5	4	3	2	1	0
0	0	0	0	0	0	0	0	0
1	Parameter Code [1]—[5]							

Notes.

[1] Parameter Code 00h.

This parameter code represents the number of logical blocks that have been sent to an initiator.

[2] Parameter Code 01h.

This parameter code represents the number of logical blocks that have been received from an initiator.

[3] Parameter Code 02h.

This parameter code represents the number of logical blocks read from the cache memory that have been sent to an initiator.

[4] Parameter Code 03h.

This parameter code represents the number of read and write commands that had data lengths equal or less than the current segment size.

[5] Parameter Code 04h.

This parameter code represents the number of read and write commands that had data lengths greater than the current segment size.

5.1.3.2.4 Factory Log Page (code 3Eh)

Table 5.1.3.2-10. Factory Log Page

Bit Byte(s)	7	6	5	4	3	2	1	0
0	0	0	0	0	0	0	0	0
1	Parameter Code [1]							

Notes.

[1] Parameter code 0000h - Power-on Time. This parameter code represents the number of drive power-on minutes. Currently the Power-on Time parameter (0000h) is the only parameter in this Log Page that is visible to OEM/customers.

5.1.3.3 Log Sense Command (4Dh)

The Log Sense command provides a means for an initiator to retrieve statistical information maintained by the drive about the drive operation. It is a complementary command to the Log Select Command. This information is stored in logs (counters) in the drive and is sent to the initiator in the data-in phase of the Log Sense command. The Log Sense command format that the initiator sends is shown in Figure 5.1.3.3-1. The format of the data pages sent back by the drive in the data-in phase is shown in Figure 5.1.3.2-3.

Note. Not supported by Standard OEM drives, but is a factory installed option.

Table 5.1.3.3-1. The Log Sense Command (4Dh)

Bit Byte(s)	7	6	5	4	3	2	1	0
0	Operation Code (4Dh)							
1	0	0	0	Reserved			PPC [1]	SP [2]
2	PC [3]		Page Code [4]					
3	Reserved							
4	Reserved							
5	(MSB)				Parameter Pointer [5]			
6	(LSB)							
7	(MSB)				Allocation Length [6]			
8	(LSB)							
9	0	0	0	0	0	0	Flag [7]	Link [7]

Notes.

[1] Parameter Pointer Control (PPC). A PPC bit of one indicates that the drive shall return a log page with only those log parameters that have changed since the last Log Select or Log Sense command. The drive returns log parameter codes according to (in ascending numerical order) the parameter code specified in the Parameter Pointer (bytes 5 & 6—see note [5]).

continued from previous page

A PPC bit of zero indicates that all of the log parameter data requested from the drive, whether changed or not, is sent and it begins with the log specified by the parameter code given in the Parameter Pointer field and returns the number of bytes specified by the allocation length field, in ascending order of parameter codes starting at the parameter code given in the Parameter Pointer field. A PPC bit of zero and a Parameter Pointer field of zero causes all available log parameters for the specified log page to be returned to the initiator subject to the allocation length.

- [2] Save Parameters bit (SP). If SP bit is zero the drive performs the specified Log Sense command and does not save any log parameters to non-volatile memory. If SP bit is one the drive first saves parameters identified as savable (by the DS bit in Table 5.1.3.2-4) to a nonvolatile location, and then performs the rest of the Log Sense command.
- [3] Page Control (PC). This field defines the type of drive log parameter the initiator requests the drive to send back on the data-in phase. See Table 5.1.3.2-2 for definition of this field. The parameter values returned on the data-in phase are from one of the following:
- (1) The specified parameter values in the log counters as of the last update (updated by Log Select command, Log Sense command or done automatically by the drive for cumulative values).
 - (2) The saved values if an update has not occurred since last power-on, hard reset condition, or Bus Device Reset message (assumes saved values are available).
 - (3) If saved values are not available, the default values are sent if an update has not occurred since the last power-on, hard reset condition, or Bus Reset message.
- [4] Page code field. The page code field specifies the page of data requested by the command (see Table 5.1.3.2-6). If a page code is sent that is not implemented, the drive terminates the command with Check Condition status. The sense key is set to Illegal Request with the additional sense code set to Invalid Field In CDB.

Table 5.1.3.3-2. (PAGE CODE 00h) returns the list of log pages that the drive supports.

Table 5.1.3.3-2. Supported Log Pages

Bit Byte(s)	7	6	5	4	3	2	1	0	
0	Reserved		Page Code (00h)						
1	Reserved								
2	(MSB)		Page Length (n-3) [1]						(LSB)
3									
4 : n	Supported Page List [2]								

This page is not defined for the LOG SELECT command.

- Notes.** [1] The page length field specifies the length in bytes of the following supported page list.
 [2] The supported page list field shall contain a list of all log page codes implemented by the target in ascending order beginning with page code 00h.

continued from previous page

- [5] **Parameter Pointer.** This field contains a Parameter Code that specifies that log parameter data be returned to the initiator by the data-in phase starting with the Parameter Pointer code log parameter data and continuing to the maximum allocation length or to (an including) log parameter data of the maximum parameter code supported by the drive, whichever is less. If the value of the Parameter Pointer field is larger than the largest available parameter code that can be returned by the drive on the specified page, 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 CDB.
- [6] **Allocation Length.** This field informs the drive of the amount of space available for returning log parameter data. The initiator can retrieve the rest of the log page information by setting the parameter pointer to the last returned parameter code and reissuing the log sense command. This process may be repeated as necessary to retrieve all the available information.
- [7] **Flag and Link bits** are used as previously defined in paragraph 4.2.6.

5.1.3.4 Other Group 2 Commands (Operation Codes 41h through 4Bh and 4Eh through 5Fh)

No other group 2 commands are implemented by the drive. A "Check Condition" status is sent if received.

5.1.4 Groups 3 through 5 Commands for All Device Types (Operation Codes 60h through BFh)

Not implemented. A "Check Condition" status is sent if received.

5.1.5 Group 6 Commands for all device types (Operation Codes C0h through DFh)

Reserved for Seagate usage. Customers should not attempt to use these functions. If these functions are used, the users stored data may be destroyed.

5.1.6 Group 7 Commands for all device types (Operation Codes E0h through FFh)

Same as Group 6 Commands.

5.2 Command descriptions for direct access devices

The drive supports Group 0, and Group 1 commands for Direct Access Devices. Refer to sections listed below for more details on these commands.

5.2.1 Group 0 Commands for Direct Access Devices

The Group 0 Commands for direct access devices implemented by drive are listed in Table 5.2.1-1, together with group 0 commands implemented for all devices.

Table 5.2.1-1. Group 0 Commands Implemented

Operation Code (HEX)	Command Name	Section	Page
00	TEST UNIT READY (All devices)	5.1.1.1	77
01	REZERO UNIT	5.2.1.1	129
03	REQUEST SENSE (All devices)	5.1.1.2	78
04	FORMAT UNIT	5.2.1.2	130
07	REASSIGN BLOCKS	5.2.1.3	137
08	READ	5.2.1.4	139
0A	WRITE	5.2.1.5	141
0B	SEEK	5.2.1.6	142
12	INQUIRY (All devices)	5.1.1.3	87
15	MODE SELECT	5.2.1.7	143
16	RESERVE	5.2.1.8	147
17	RELEASE	5.2.1.9	149
1A	MODE SENSE	5.2.1.10	150
1B	START/STOP UNIT	5.2.1.11	183
1C	RECEIVE DIAGNOSTIC RESULTS (All devices)	5.1.1.5	100
1D	SEND DIAGNOSTIC (All devices)	5.1.1.6	105

5.2.1.1 Rezero Unit Command (01h)

Table 5.2.1-2. Rezero Unit command

Bit Byte(s)	7	6	5	4	3	2	1	0
0	0	0	0	0	0	0	0	1
1	Logical Unit No. 0 0 0 [1]			0	0	0	0	0
2	0	0	0	0	0	0	0	0
3	0	0	0	0	0	0	0	0
4	0	0	0	0	0	0	0	0
5	0	0	0	0	0	0	Flag [2]	Link [2]

The Rezero Unit command (Table 5.2.1-2) requests that the drive set its logical block address to zero and return the drive read/write heads to the track (or cylinder) containing Logical Block Zero. This command is intended for systems which disable retries and the initiator performs error recovery. It is longer than a seek to Logical Block Address zero and should be utilized if seek errors are encountered.

For systems that support disconnection, the drive disconnects when this command is received.

A Rezero Command also causes a thermal compensation to occur and resets the thermal compensation cycle timer back to its start, thus allowing the host to know when to expect the next thermal compensation to occur. The host can thus prevent critical data transfer operations from being interrupted at an undesirable time.

For drives that support saved log parameters, the Rezero Unit command will also save log counters to the media and reset the log save timer back to its start.

Notes.

[1] The LUN must be zero.

[2] See "Control Byte" paragraph 4.2.6.

5.2.1.2 Format Unit Command (04h)

Table 5.2.1-3. Format Unit Command

Bit Byte(s)	7	6	5	4	3	2	1	0
0	0	0	0	0	0	1	0	0
1	Logical Unit No. [1] 0 0		0	FMT [2] DATA	CMP [2] LST	Defect List Format [2]		
2	0	0	0	0	0	0	0	0
3 4	INTERLEAVE [3]							
5	0	0	0	0	0	0	Flag [4]	Link [4]

The Format Unit command (Table 5.2.1-3) ensures that the medium is formatted so all of the user addressable data blocks can be accessed. There is no guarantee that the medium has or has not been altered. In addition, the medium may be certified and control structures may be created for the management of the medium and defects.

For systems which support disconnection, the drive disconnects while executing the Format Unit command.

This command is implemented in the drive for mandatory features and a subset of the available optional features of the ANSI SCSI-2 specification as defined in the following paragraphs.

The drive allows an initiator to specify (or not specify) sectors which are to be reallocated during the format process. The format parameters to be specified in the Format Unit command are defined in Table 5.2.1-4.

The FORMAT UNIT command shall be rejected with RESERVATION CONFLICT status if the logical unit is reserved, or any extent reservation, from any initiator, is active in the specified logical unit.

Notes.

[1] The LUN must be zero.

[2] A Format Data (FmtData) bit of one indicates that the Format Unit Parameter list (Table 5.2.1-5a) is supplied during the Data Out phase. The Data Out phase consists of a defect list header (Table 5.2.1-5b), followed by an initialization pattern descriptor (Table 5.2.1-6a) (if any) followed by the defect descriptors. The format of the defect descriptor list is determined by the Defect List Format field. A FmtData bit of zero indicates the Data Out phase shall not occur (no defect data shall be supplied by the initiator).

A Complete List (CmpLst) bit of one indicates the data supplied is to be the complete list of Growth defects. Any previous Growth defect data or Certification defect data shall be erased. The drive may add to this list as it formats the medium. The result is to purge any previous Growth or Certification defect list and to build a new defect list. A CmpLst bit of zero indicates the data supplied is in addition to the existing Growth defect list.

The use of the P list and C list defect is controlled by byte 1 of the defect list header (see Table 5.2.1-5b).

The Defect List Format field specifies additional information related to the defect list. (See Table 5.2.1-4 for further information.)

[3] The Interleave field requests that logical blocks be related in a specific fashion to the physical blocks to facilitate data transfer speed matching. An interleave value of zero requests that the target use its default interleave. An interleave value of one requests that consecutive logical blocks be placed in consecutive physical order. Values of two or greater indicate that one or more (respectively) physical blocks separate consecutive logical blocks. The drive implements an optional prefetch ("read look ahead") function which reads a user-specified number of sectors into its buffer beyond and contiguous to the sectors requested by the read command. This data is subsequently available for the next sequential read command without re-access of the disc media thereby increasing performance and negating the need for an interleave during format if this prefetch option is enabled. (See Read Command, Section 5.2.1.4.) Interleave values other than zero or one are vendor specific.

[4] See "Control Byte" paragraph 4.2.6.

The following definitions of flaw categories are supplied to help in understanding the alternatives listed in Table 5.2.1-4.

P = Primary Defect Type: P type flawed sectors are identified at the time of shipment in a list of defects (permanent flaws) supplied by Seagate and stored on the disc in an area that is not directly accessible by the user. (This list may be referred to as an ETF List). This defect list is not modified or changed by the drive (or initiator) after shipment.

C = Certification Defect Type: C type flawed sectors are sectors that fail a format verify during the format function.

D = Data Defect Type: D type sectors are sectors identified in a list supplied to the target by the initiator during a Data Out phase of the current Format Unit command. The D List follows a four byte defect list header and is referred to as Defect Descriptor Bytes.

G = Growth Defect Type: G type flawed sectors contain medium flaws and have been reallocated as a result of receiving a Reassign Blocks command, or certification defects (C type) reallocated during a previous Format Unit command, or Data Defects (D type) reallocated during a previous Format Unit command or defects that have been automatically reallocated by the drive. This (G) list is recorded on the drive media and may be referenced for the current (and subsequent) Format Unit commands. This (G) list does not include the Primary (P) list of defects.

Table 5.2.1-4. Format Unit Parameter Definition (Format Variations)

CDB-Byte [1]					COMMENTS
Bits					
4	3	2	1	0	
FMT DATA	CMP LIST	DEFECT LIST FORMAT			
0	X	X	X	X	Default Format: No Data Out phase occurs. Drive reallocates all sectors in the P list plus any sector which fails the Format Verify phase (C type flaws). Any previous G list is erased.
Block format [3]					
1	0	0	X	X	Format with G and no D: A four byte Defect List Header must be sent by the initiator. No Defect Descriptors (D list) are sent by the initiator. The drive reallocates all sectors in the drives current G list. See also note [2].
1	1	0	X	X	Format without G or D: A four byte Defect List Header must be sent by the initiator. No D list may be sent by the initiator. The drive erases any previous G list. See also note [2].
Bytes from INDEX format					
1	1	1	0	0	Format with D and without G. The initiator must send a four byte Defect List Header followed by a D list of the defects for the drive to reallocate. The D list must be in the bytes from Index format (see Table 5.2.1-6e). The drive erases any previous G list. See also Note [2].
1	0	1	0	0	Format with D and with G: The initiator must send a four byte Defect List Header followed by a D list of the defects for the drive to reallocate. The D list must be in the bytes from Index format (see Table 5.2.1-6e). The drive also reallocates all sectors in the drives current G list. See also note [2].
Physical Sector format					
1	1	1	0	1	Format with D and without G: The initiator must send a four byte Defect List Header followed by a D List of defects for the drive to reallocate. The D list must be in the Physical Sector format (see Table 5.2.1-6f). The drive erases any previous G list. See also note [2].
1	0	1	0	1	Format with D and with G: The initiator must send a four byte Defect List Header followed by a D List of defects for drive to reallocate. The D list must be in the Physical Sector format (see Table 5.2.1-6f). The drive also reallocates all sectors in the drives current G list. See also note [2].

[1] Refer to Table 5.2.1-3.

[2] Byte one of the Defect List Header determines whether the P and C defects are reallocated. See Table 5.2.1-5b.

[3] See individual Volume 1 Product Manuals for support/non-support.

The defect list shown in Table 5.2.1-5b contains a four byte header followed by one or more defect descriptors (Table 5.2.1-6d). The Defect List Length in each table specifies the total length in bytes of the defect descriptors that follow. In Table 5.2.1-5b the Defect List Length is equal to eight times the number of defect descriptors.

Table 5.2.1-5a. Format Unit Parameter List

Bit Byte(s)	7	6	5	4	3	2	1	0
*	DEFECT LIST HEADER				* (Table 5.2.1-5b)			
**	INITIALIZATION PATTERN DESCRIPTOR				** (Table 5.2.1-6a)			
***	DEFECT DESCRIPTOR(S)				*** (Tables 5.2.1-6d, 6e, 6f)			

Table 5.2.1-5b. Defect List Header

Bit Byte(s)	7	6	5	4	3	2	1	0
0	RESERVED							
1	FOV [1]	DPRY [2]	DCRT [3]	STPF [4]	IP [5]	DSP [6]	IMMED [7]	VS [8]
2	(MSB) DEFECT LIST LENGTH [9]							(LSB)
3								

Notes.

Function Drive defect List Header Bit Interpretations

[1] FOV If one, the drive interprets the remaining bits of byte 1. If zero, the drive checks the remaining bits of byte 1 for zeros.

[2] DPRY If one, flaws in the drive P list are not reallocated during formatting. This means existing reallocations of the P list are canceled and no new reallocations made during formatting. The P list is retained. Some Seagate drives do not support a DPRY bit of one. See individual drive Product Manual Volume 1.

If zero, flaws in the drive P list are reallocated during formatting. A Check Condition is sent in the status if the P list cannot be found by the drive.

[3] DCRT If one, the drive does not perform a verify function during formatting (thus no C list for this format is created or reallocated).

If zero, the drive performs a verify function during formatting and reallocates any sector that fails the verify (i.e.; a C list is created and these flaws reallocated).

On drives that have MR heads the DCRT_BIT is always interpreted as a 0 when the user sets the IP bit (see Table 5.2.1-5b) and specifies a format pattern of **greater** than 1 byte (see Table 5.2.1-6a), and/or if the IP MODIFIER bits are set to 01 or 10 (see Table 5.2.1-6a). This is done since the user format pattern and/or IP header is written onto the media during the format certification phase. Because of this the user may see G list entries added to the defect list even though the DCRT bit was set to 1.

[4] STPF If one, formatting is terminated with a CHECK CONDITION STATUS if an error is encountered while accessing either the P or G defect list. The Sense key is set to MEDIUM ERROR and the additional sense code shall be set to either DEFECT LIST NOT FOUND or DEFECT LIST ERROR.

If zero, formatting is not terminated if an error is encountered while accessing either the P or G defect list.

continued from previous page

- [5] IP An initialization pattern (IP) bit of one indicates that an initialization pattern descriptor is included in the FORMAT UNIT parameter list immediately following the defect list header. An IP bit of zero indicates that an initialization pattern descriptor is not included and that the target uses its default initialization pattern.
- The initialization pattern type field indicates the type of pattern the target shall use to initialize each logical block within the initiator accessible portion of the medium (Table 5.2.1-6c). All bytes within a logical block shall be written with the initialization pattern. The initialization pattern is modified by the IP modifier field as described in Table 5.2.1-6b.
- [6] DISABLE SAVING PARAMETERS A disable saving parameters (DSP) bit of one, specifies that the target shall not save the MODE SELECT savable parameters to non-volatile memory during the format operation. A DSP bit of zero specifies that the target shall save all the MODE SELECT savable parameters for all initiators to non-volatile memory during the format operation.
- [7] IMMEDIATE An immediate (IMMED) bit of zero indicates that status shall be returned after the first operation has completed. An Immed bit value of one indicates that the target shall return status as soon as the command descriptor block has been validated, and the entire defect list has been transferred. Check individual drive Product manual Vol. 1 for applicability of this feature.
- [8] VS Vendor Specific. Not used.
- [9] DEFECT LIST LENGTH The length of any following D list (Defect Descriptors) must be equal to 8 times the number of sectors to be reallocated per Table 5.2.1-6e or Table 5.2.1-6f.

Note. The initiator may not use any previously defined “C”, “G”, or “D” lists if the sector size (block length) has been changed.

The initialization pattern option specifies that the logical blocks contain the specified initialization pattern. The initialization pattern descriptor (Table 5.2.1-6a) is sent to the target as part of the FORMAT UNIT parameter list.

Table 5.2.1-6a. Initialization Pattern Descriptor

Bit Byte(s)	7	6	5	4	3	2	1	0
0	IP MODIFIER [1]		RESERVED					
1	PATTERN TYPE							
2	(MSB)		PATTERN LENGTH [2]					
3								
0-n	INITIALIZATION PATTERN [3]							

Table 5.2.1-6b. Initialization Pattern Modifier

This table gives pattern type information required in byte 1 of Table 5.2.1-61.

IP	Modifier [1]	Description
0	0	No header. The target shall not modify the initialization pattern.
0	1	The target shall overwrite the initialization pattern to write the logical block address in the first four bytes of the logical block. The logical block address shall be written with the most significant byte first.
1	0	The target shall overwrite the initialization pattern to write the logical block address in the first four bytes of each physical block contained within the logical block. The lowest numbered logical block or part thereof that occurs within the physical block is used. The logical block address shall be written with the most significant byte first.
1	1	Reserved.

Table 5.2.1-6c. Initialization Pattern Type

Pattern Type	Note	Description
00h	[4]	Use default pattern
01h	[5]	Use pattern supplied by host computer
02-7Fh	None	Reserved
80-FFh	None	Vendor-specific

Notes. (These notes apply to Tables 5.2.1-6a, 5.2.1-6b and 5.2.1-6c)

- [1] The initialization pattern [3] is modified in accordance with the specification of the IP MODIFIER field given in Table 5.2.1-6a.
- [2] The initialization pattern length field indicates the number of bytes contained in the initialization pattern [3]. If the length exceeds the current logical block size this is an error. See Note [6].
- [3] The initialization pattern is sent by the host for the drive to write in each logical block by the format command.
- [4] If the initialization pattern length [2] is not zero, this is an error. See Note [6].
- [5] If the initialization pattern length [2] is zero, this is an error. See Note [6].
- [6] If an error occurs in Notes [2], [4] or [5] the drive shall terminate the command with CHECK CONDITION status. The sense key shall be set to ILLEGAL REQUEST and the additional sense error code shall be set to INVALID FIELD IN PARAMETER LIST.

Defect List Formats

Table 5.2.1-6d. Defect Descriptors

0-n	DEFECT DESCRIPTORS BYTES (Refer to Table 5.2.1-6e, Table 5.2.1-6f for Format)
-----	--

This section describes the format of the defect list that follows the Defect List Header described in Table 5.2.1-5b. Three formats are possible: the block format, the bytes from index format and the physical sector format. Seagate drives do not support the block format, except as a customer special. The other two formats that are supported are described in Tables 5.2.1-6e and 5.2.1-6f.

Table 5.2.1-6e. Defect Descriptor Bytes - Bytes from Index Format

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

For defects to be specified in the Bytes from Index format, the defect list format field (Byte 1, bits 2, 1, 0) must be 100 (binary), see Tables 5.2.1-3 and 5.2.1-4.

Each Defect Descriptor for the Bytes from Index format specifies the beginning of an eight bit (1 byte) defect location on the medium. Each defect descriptor is comprised of the cylinder number of the defect, the head number of the defect and the number of bytes from Index to the defect location. (Defect bytes from Index)

The Defect Descriptors shall be in ascending order. The drive may return CHECK CONDITION if the defect descriptors are not in ascending order. For determining ascending order, the Cylinder Number of Defect is considered the most significant part of the address and the Defect Bytes from Index is considered the least significant part of the address.

A value for Defect Bytes from Index of FFFFFFFFh (i.e., reassign the entire track) is illegal for the drive.

Table 5.2.1-6f. Defect Descriptor Bytes - Physical Sector Format

Byte Numbers	Description
0	Cylinder Number of Defect (MSB)
1	Cylinder Number of Defect
2	Cylinder Number of Defect (LSB)
3	Head Number of Defect
4	Defect Sector Number (MSB)
5	Defect Sector Number
6	Defect Sector Number
7	Defect Sector Number (LSB)

Information in this Table is repeated for each defect.

For defects to be specified in the Physical Sector format, the Defect List format field (Byte 1, bits 2, 1, 0) must be 101 (binary). See Tables 5.2.1-3 and 5.2.1-4.

Each Defect Descriptor for the Physical Sector format specifies a sector size defect location comprised of the cylinder number of the defect, the head number of the defect and the defect sector number.

The Defect Descriptors shall be in ascending order. The drive may return CHECK CONDITION if the defect descriptors are not in ascending order. For determining ascending order, the Cylinder Number of the defect is considered the most significant part of the address and the Defect Sector Number is considered the least significant part of the address.

A Defect Sector Number of FFFFFFFFh (i.e., reassign the entire track) is illegal for the drive.

5.2.1.3 Reassign Blocks Command (07h)

Table 5.2.1-7. Reassign Blocks Command

Bit Byte(s)	7	6	5	4	3	2	1	0
0	0	0	0	0	0	1	1	1
1	Logical Unit No.[1] 0 0 0			0	0	0	0	0
2	0	0	0	0	0	0	0	0
3	0	0	0	0	0	0	0	0
4	0	0	0	0	0	0	0	0
5	0	0	0	0	0	0	Flag	Link [2]

Notes.

[1] The LUN must be zero.

[2] See "Control Byte" paragraph 4.2.6.

The Reassign Blocks command (Table 5.2.1-7) requests the target to reassign the defective logical blocks to an area on the logical unit reserved for this purpose.

After sending the Reassign Blocks command, the initiator transfers a defect list that contains the logical block addresses to be reassigned. The drive reassigns the physical medium used for each logical block address in the list. The data contained in the logical blocks specified in the defect list is not preserved, but the data in all other logical blocks on the medium is preserved. It is recommended that the initiator recover the data from the logical blocks to be reassigned before issuing this command. After completion of this command, the initiator can write the recovered data to the same Logical Block Addresses.

The effect of specifying a logical block to be reassigned that has previously been reassigned is to reassign the block again. Thus, over the life of the medium, a logical block can be assigned to multiple physical addresses (until no more spare locations remain on the medium).

This command should be used by an initiator to immediately reallocate any block (sector) which requires the drive to recover data by data correction via ECC if the automatic reallocation feature of the drive is not enabled, see Mode Select command (Section 5.2.1.7).

For systems which support disconnection, the drive disconnects while executing this command.

The Reassign Blocks defect list (Table 5.2.1-8) contains a four byte header followed by one or more Defect Descriptors. The length of each Defect Descriptor is four bytes.

Table 5.2.1-8. Reassign Blocks Defect List

Defect List Header								
Bit Byte(s)	7	6	5	4	3	2	1	0
0	0	0	0	0	0	0	0	0
1	0	0	0	0	0	0	0	0
2	Defect List Length (MSB)						[1]	
3	Defect List Length (LSB)							
Defect Descriptor(s)								
0	Defect Logical Block Address (MSB)							[2]
1	Defect Logical Block Address							
2	Defect Logical Block Address							
3	Defect Logical Block Address (LSB)							

Notes.

[1] The Defect List Length specifies the total length in bytes of the Defect Descriptors that follow. The Defect List Length is equal to four times the number of Defect Descriptors.

[2] The Defect Descriptor specifies a four byte Defect Logical Block Address that contains the defect. The Defect Descriptors shall be in ascending order.

If the logical unit has insufficient capacity to reassign all of the defective logical blocks, the command shall terminate with a Check Condition status and the Sense Key shall be set to HARDWARE ERROR and the additional sense code set to NO DEFECT SPARE LOCATION AVAILABLE. The logical block address of the first logical block not reassigned shall be returned in the information bytes of the sense data. If information about the first defect descriptor not reassigned is not available, or if all the defects have been reassigned, this field shall be set to FFFFFFFFh.

If the REASSIGN BLOCKS command failed due to an unexpected unrecoverable read error that would cause the loss of data in a block not specified in the defect list, the logical block address of the unrecoverable block shall be returned in the information field of the sense data and the valid bit shall be set to one.

IMPLEMENTORS NOTE: If the REASSIGN BLOCKS command returns CHECK CONDITION status and the sense data command-specific information field contains a valid logical block address, the initiator should remove all defect descriptors from the defect list prior to the one returned in the command-specific information field. If the sense key is MEDIUM ERROR and the valid bit is one (the information field contains the valid block address) the initiator should insert that new defective logical block address into the defect list and reissue the REASSIGN BLOCKS command with the new defect list. Otherwise, the initiator should perform any corrective action indicated by the sense data and then reissue the REASSIGN BLOCKS command with the new defect list.

5.2.1.4 Read Command (08h)

Table 5.2.1-9. Read Command

Bit Byte(s)	7	6	5	4	3	2	1	0
0	0	0	0	0	1	0	0	0
1	Logical Unit No.[1] 0 0 0			Logical Block Address (MSB) [2]				
2	Logical Block Address							
3	Logical Block Address (LSB)							
4	Transfer Length [3]							
5	0	0	0	0	0	0	Flag	Link [4]

The Read Command (Table 5.2.1-9) requests that the drive transfer data to the initiator.

Notes.

[1] In the CDB the LUN must be zero.

[2] The Logical Block Address specifies that logical block at which the read operation shall begin.

[3] The Transfer Length specifies the number of contiguous logical blocks of data to be transferred. A Transfer Length of zero indicates that 256 logical blocks shall be transferred. Any other value indicates the number of logical blocks that shall be transferred.

[4] See "Control Byte" paragraph 4.2.6.

The data value most recently written in the addressed logical block shall be returned.

Read data transfers with the initiator do not begin until at least one full sector of data is available in the drive data buffer. For multiple sector reads, the transfer of data continues until the number of blocks specified in byte 4 of the CDB has been read and transferred or until an unrecoverable error is detected.

Data transfer could stop if the option to stop on recovered error is selected.

For systems that support disconnection, the drive disconnects when a valid Read command is received. The drive reconnects depending on the value of the Buffer Full Ratio Set in Page 2 of the Mode Select Data (see Section 5.2.1.7). After data transfer has been initiated with an initiator, the drive does not disconnect unless an internal error recovery procedure is required or the data transfer to an initiator is interrupted for more than 1 millisecond.

The initiator must accept all data presented to the initiator after sending this command until the drive sends Completion Status during a Status phase. (Note the drive may disconnect and reconnect while executing this command and the initiator may prematurely terminate this command by creating the Reset condition or by sending an Abort, Abort Tag, Clear Queue or Bus Device Reset message).

Sense Data is valid after this command is executed and Completion Status is sent. If the Address Valid bit in the Sense Data is true (1), the Sense Data Logical Block Address (Information bytes) points to the last Logical Block accessed by the drive. If the Address Valid bit in the Sense Data is false (0), the Sense Data Logical Block Address bytes are not valid.

The drive contains a large buffer and implements an optional “prefetch” and segmented cache function whereby the requested read data is read into the buffer, plus an additional amount, depending on the cache control parameters. See paragraph “Prefetch and multisegmented cache control “ in the individual Product Manual for more information on this.

This command is terminated with a Reservation Conflict status and no data is read if any reservation access conflict (see Section 5.2.1.8) exists.

If any of the following conditions occur, this command shall be terminated with a Check Condition status, and if extended sense is implemented, the Sense Key shall be set as indicated in the following table. This table does not provide an exhaustive enumeration of all conditions that may cause the Check Condition status.

Condition	Sense Key
Invalid Logical Block Address	Illegal Request (see note)
Target reset since last command from this initiator	Unit Attention
Unrecoverable read error	Medium Error
Recovered read error	Recovered Error

Note. The extended sense Information Bytes shall be set to the Logical Block Address of the first invalid address.

5.2.1.5 Write Command (0Ah)

Table 5.2.1-10. Write Command

Bit Byte(s)	7	6	5	4	3	2	1	0
0	0	0	0	0	1	0	1	0
1	Logical Unit No. [1] 0 0 0			Logical Block Address (MSB) [2]				
2	Logical Block Address							
3	Logical Block Address (LSB)							
4	Transfer Length							[3]
5	0	0	0	0	0	0	Flag	Link [4]

The Write command (Table 5.2.1-10) requests that the drive write, to the medium, the data transferred by the initiator.

Notes.

[1] The LUN must be zero.

[2] The Logical Block Address specifies the logical block at which the write operation shall begin.

[3] The Transfer Length specifies the number of contiguous logical blocks of data to be transferred. A Transfer Length of zero indicates that 256 logical blocks shall be transferred. Any other value indicates the number of logical blocks that shall be transferred.

[4] See "Control Byte" paragraph 4.2.6.

For systems that support disconnection, the drive disconnects when any internal error recovery procedure is required, or the data transfer with the initiator is interrupted for more than 1 millisecond, or if the drive's internal data buffer is full. After a disconnect the drive reconnects depending on the value of the Buffer Empty Ratio in Page 2 of Mode Select Data (see Section 5.2.1.7).

The initiator must send requested write data to the drive until the drive sends Completion status during a Status phase or until the initiator Resets/Aborts the command or clears the queue. (Note: the drive may disconnect and reconnect while executing this command).

Sense Data is valid after this command is executed and Completion status is sent. (refer to the Read Command description in Section 5.2.1.4).

If the RCD bit is set to zero on the Caching Mode page 08h (cache is enabled), the data that is written by this command remains in the cache buffer, if no write errors are encountered. This allows a Read command to access the same data from the cache buffer instead of accessing the media, if the same LBA is requested by the Read command.

This command is terminated with a Reservation Conflict status and no data is written if any reservation access conflict (see Section 5.2.1.8) exists.

If any of the following conditions occur, this command is terminated with a Check Condition status, and if extended sense is implemented, the Sense Key is set as indicated in the following table. This table does not provide an exhaustive enumeration of all conditions that may cause the Check Condition status.

Condition	Sense Key
Invalid Logical Block Address	Volume overflow (see note)
Target reset since last command from this initiator	Unit Attention

Note. The extended sense Information Bytes shall be set to the Logical Block Address of the first invalid address.

5.2.1.6 Seek Command (0Bh)

Table 5.2.1-11. Seek Command

Bit Byte(s)	7	6	5	4	3	2	1	0
0	0	0	0	0	1	0	1	1
1	Logical Unit No.[1] 0 0 0			Logical Block Address (MSB) [2]				
2	Logical Block Address							
3	Logical Block Address (LSB)							
4	0	0	0	0	0	0	0	0
5	0	0	0	0	0	0	Flag	Link [3]

The Seek command (Table 5.2.1-11) requests that the drive seek to the specified logical block address.

For systems which support disconnection, the drive disconnects when a valid Seek command is received. The use of this command is infrequent since all commands involving data transfer to/from the drive media contain implied seek addresses.

Notes.

[1] The Logical Unit number should be zero.

[2] The maximum Logical Block Address that may be specified for a Seek command is defined in Read Capacity Data, Section 5.2.2.1.

[3] See "Control Byte" paragraph 4.2.6.

5.2.1.7 Mode Select (6) Command (15h)

Table 5.2.1-12. Mode Select (6) Command

Bit Byte(s)	7	6	5	4	3	2	1	0
0	0	0	0	1	0	1	0	1
1	Logical Unit No.[1] 0 0 0			PF [2]	0	0	0	SMP [3]
2	0	0	0	0	0	0	0	0
3	0	0	0	0	0	0	0	0
4	Parameter List Length							[4]
5	0	0	0	0	0	0	Flag	Link [5]

The Mode Select command (Table 5.2.1-12) provides a means for the initiator to specify medium, logical unit, or peripheral device parameters to the drive. The drive also implements the Mode Sense command (See 5.2.1.10). Initiators should issue Mode Sense prior to Mode Select to determine supported pages, page lengths, and other parameters.

The drive maintains a separate set of mode parameters for each initiator that could be on the SCSI bus (7 or 15), regardless of whether or not more than one initiator actually exists on the bus. If an initiator sends a Mode Select command that changes any parameters that apply to other initiators, the drive generates a Unit Attention condition for all initiators except the one that issued the Mode Select command. The drive sets the additional sense code to Mode Parameters Changed.

Notes.

[1] LUN must be zero

[2] The Page Format (PF) bit, when set to one, indicates the data sent by the initiator after the mode select Header and Block Descriptors (if any) complies with the Page Format. The PF bit, when set to zero, indicates the data sent after the Mode Select Header and the Block Descriptors (if any) are vendor unique. The drive does not interpret the PF bit. It assumes the Page Format mode.

continued from previous page

- [3] The Save Mode Parameters (SMP) bit, when set to one, requests that the drive save the savable pages. The format related parameters in the block descriptor, pages 3 and 4 are saved during a Format command as well as a Mode Select command with SMP = 1. The drive must update the Current mode values with parameters included with this command, save the Current values of the savable parameters, and report Good status only after the save operation is completed. The Saved parameters are not changed if an error is detected during the Mode Select command. When the SMP bit is set to zero, the Saved parameter values are not changed.
- [4] The Parameter List Length specifies the length in bytes of the Mode Select parameter list that shall be transferred during the Data Out phase. A Parameter List Length of zero indicates that no data shall be transferred. This condition shall not be considered as an error.
- [5] See "Control Byte" paragraph 4.2.6.

The Mode Select parameter list (Table 5.2.1-13) contains a four byte header, followed by zero or one block descriptor, followed by the pages of Mode Select Parameters.

Acceptable values for the Mode Select parameter list for the drive are shown in Table 5.2.1-13.

The target terminates all the Mode Select commands with Check Condition status, sets the sense key to ILLEGAL REQUEST and sets the additional sense code to Invalid Field In Parameter List, *and does not change any mode parameters* for the following conditions:

1. If *the Strict mode is enabled (see note [5] of Table 5.2.1-31) and the initiator attempts to change any field that is not changeable by the host as reported by the target. In this case, no parameters are changed by this command. The target compares the parameters against the values as they were prior to this Mode Select command. (The host shall not be penalized by values, not changeable by the host, which have a target "ripple change" as a result of this Mode Select).*
2. If the initiator attempts to send an unsupported value or, a non-zero value to a reserved field in the Mode Select header, block descriptor, or any page header.
3. If an initiator attempts to send a page with a length not equal to the parameter length reported for that page by the Mode Sense command.
4. If the initiator attempts to send a value for a *changeable* parameter that is outside the range supported by the target and rounding is not implemented for that parameter (see note [4], Table 5.2.1-31).
5. If the initiator sends a page descriptor with an unsupported page code value and the Strict mode is enabled. See note [5] of Table 5.2.1-31.

If the initiator sends a value for a *changeable* parameter that is outside the range supported by the target and rounding is implemented for that parameter, the target shall either:

- (1) round the parameter to an acceptable value and *if Round is one*, terminate the command as described in 4.8.
- (2) *round the parameter to an acceptable value and if Round equals zero, terminate the command as if an acceptable value had been sent from the initiator.*

A target may alter any mode parameter in any mode page (even parameters reported as non-changeable) as a result of changes to other mode parameters.

Table 5.2.1-13. Mode Select Parameter List

Bit Byte(s)	7	6	5	4	3	2	1	0
0	0	0	0	0	0	0	0	0
	RESERVED							
1	0	0	0	0	0	0	0	0
	MEDIUM TYPE [1]							
2	0	0	0	0	0	0	0	0
	RESERVED							
3	BLOCK DESCRIPTOR LENGTH EITHER 0 or 8 (DECIMAL) [2]							
	BLOCK DESCRIPTOR							
0	DENSITY CODE [3]							
1	NUMBER OF BLOCKS (MSB) [3] [4]							
2	NUMBER OF BLOCKS [4]							
3	NUMBER OF BLOCKS (LSB) [4]							
4	0	0	0	0	0	0	0	0
5	BLOCK LENGTH (MSB)							
6	BLOCK LENGTH [5]							
7	BLOCK LENGTH (LSB)							
	PARAMETER INFORMATION [6]							
0 - n	MODE SELECT PAGE HDRS AND THEIR PARAMETERS (Tables 5.2.1-14 and 5.2.1-15)							

Notes.

- [1] The medium type field shall be 00h to define the default type direct access device.
- [2] The Block Descriptor Length specifies the length in bytes of the Block Descriptor. It is equal to the number of bytes in the Block Descriptor (either 0 or 8) and does not include the page headers and mode parameters. A Block Descriptor Length of zero indicates that no block descriptors shall be included in the parameter list. This condition shall not be considered an error.
- [3] a. For drives that do not support capacity programming, the Density Code shall be 00h to define the default density of medium.
b. If the drive supports capacity programming (see note [4]), byte 0 is useable as MSB part of the number in bytes 1, 2 and 3. For drives whose capacity does not require the use of byte 0, byte 0 will always be 00h.
- [4] a. For drives that do not support capacity programming, these bytes are always zero.
b. A value of zero in bytes 1, 2 and 3 indicates that the drive shall not change the capacity it is currently formatted to have. Any other value in these bytes is ignored by drives not having the capacity programming feature. For drives that have the capacity programming capability (see Product Manual Volume 1), a number in bytes 0, 1, 2 and 3 that is less than the maximum number of LBA's changes the drive capacity to the value in the block descriptor bytes 0, 1, 2 and 3. A value greater than the maximum number of LBA's is rounded down to the maximum capacity.
- [5] Block Length specifies the length in number of bytes for each logical block described by the Block Descriptor. Set to desired sector size before a Format. (Valid values are even numbered sizes from 180 to 4096. Not all drives can format down to 180; some have a minimum of 256).
- [6] See Mode Sense Command, for detailed descriptions of the Mode Select/Sense pages.

The rest of the Mode Select parameters are organized into pages that group the parameters by function. The parameter definitions are the same as those described in the Mode Sense command (paragraph 5.2.1.10) and are not repeated here.

Table 5.2.1-14. Mode Select Page Descriptor Header

Bit Byte(s)	7	6	5	4	3	2	1	0
0	0	0	PAGE CODE					
1	PAGE LENGTH							
2 - n	MODE PARAMETERS							

Each page of mode parameters begins with a two byte Page Descriptor Header. The Page Code identifies the page of mode parameters that is being transferred. The Page Length indicates the number of additional bytes of mode parameters contained in this page. The number of additional bytes sent must always match the Page Length value.

*The drive only verifies Mode Select Data that is defined as changeable by the drive.** The various drives support the following Page Codes*:

Table 5.2.1-15. Mode Sense Page Codes supported

Page Code	Description
00h	Unit Attention Page parameters (should be sent last in a group of Mode pages)
01h	Error Recovery parameters
02h	Disconnect/Reconnect Control parameters
03h	Format parameters
04h	Rigid Drive Geometry parameters
07h	Verify Error Recovery parameters
08h	Caching Parameters
0Ah	Control Mode page
0Ch or 1A	Notch and Partition page
0Dh	Power Condition page
10h	Xor Control Mode page
1Ch	Information Exceptions Control page
3Fh	Return all supported pages

The detailed information can be obtained by issuing the Mode Sense command requesting changeable values.

Note.

There may be implicit associations between parameters defined in the pages and block descriptors. The block length affects the optimum values (the values that achieve best performance) for the sectors per track, bytes per physical sector, track skew factor, and cylinder skew factor fields in the format parameters page. In this case, the drive may change parameters not explicitly sent with the Mode Select command. A subsequent Mode Sense command would provide information on these changes.

* See individual drive Product Manuals (Volume 1), Section "SCSI Interface commands supported", for a table showing the mode pages that a particular drive implements. The table shows the default parameters for pages that are implemented, and shows which mode parameters are changeable by that drive model.

5.2.1.8 Reserve Command (16h)

The Reserve and Release commands provide the basic mechanism for contention resolution in multiple initiator systems. The third party reservation allows one initiator to reserve the drive for some other SCSI device that can act as an initiator. Table 5.2.1-16 shows the command structure.

Table 5.2.1-16. Reserve Command CDB

Bit Byte(s)	7	6	5	4	3	2	1	0
0	0	0	0	1	0	1	1	0
1	Logical Unit No. [1]		3rd PRTY 0 or 1 [2]	3rd PRTY DEV.ID [2]			EXTENT [3]	
2	RESERVATION IDENTIFICATIONS							
	0	0	0	0	0	0	0	0 [4]
3	EXTENT LIST LENGTH (MSB)							
	0	0	0	0	0	0	0	0 [3]
4	EXTENT LIST LENGTH (LSB)							
	0	0	0	0	0	0	0	0 [3]
5	0	0	0	0	0	0	Flag [5]	Link [5]

Notes.

[1] The LUN must be zero.

[2] If bit 4 is zero, bits 3, 2, and 1 are zeros. If bit 4 is one, bits 3, 2, and 1 identify the SCSI bus ID of the device for which the drive is reserved.

[3] Must be zero if not supported. Check with drive Product Manual, "SCSI Interface commands supported".

[4] Reserve Identifications applies only if Extent is supported. This field must be zero if not supported. Check with drive Product Manual, section "SCSI Interface commands supported". (also see paragraph 5.2.1.8.1, 5.2.1.8.2).

[5] See "Control Byte" paragraph 4.2.6.

5.2.1.8.1 Logical Unit Reservation

If the Extent bit is zero, this command requests the entire drive be reserved for exclusive use of the initiator until the reservation is superseded by another valid Reserve command from the initiator that made the reservation, released by a Release command from the same initiator, by a Bus Device Reset message from any initiator, or by a "hard" Reset condition. A logical unit reservation is not granted if any extent or logical unit is reserved by another initiator or if any extent with a read shared reservation type is reserved by this initiator. It is permissible for an initiator to reserve a logical unit that is currently reserved by that initiator. If the Extent bit is zero, the Reservation Identification and the Extent List Length are ignored. If the Extent bit is a one and if the drive does not support the extent reservation option, the drive generates Check Condition status and sets the sense key to Illegal Request.

If, after honoring the reservation, any other initiator subsequently attempts to perform a command other than a Release command, which shall be ignored, or an Inquiry command, which shall be executed, or a Request Sense command, which shall be executed, the command is rejected with Reservation Conflict status.

5.2.1.8.2 Extent Reservation

The drive may not support Extent reservations, in which case this bit must always be zero. See the drive Product Manual section “SCSI Interface commands supported”. Since the Reservation Identification byte and the Extent List Length are valid only for extent reservations, the drive ignores these fields if extent reservations are not supported.

5.2.1.8.3 Third Party Reservation

The third party reservation option for the Reserve command allows an initiator to reserve a logical unit for another SCSI device. If the drive supports third party reserve command, the drive also implements the third party release option (see Section 5.2.1.9.3). This feature is intended for use in Multiple-initiator systems that use the Copy command. Check the Drive Product Manual to see if the drive supports the third party reservation option.

If the third party (3rdPty) bit is zero, the third party reservation option is not requested. If the 3rdPty bit is one the Reserve command reserves the specified logical unit for the SCSI device specified in the third party device ID field. The drive shall preserve the reservation until it is superseded by another valid Reserve command from the initiator which made the reservation or until it is released by the same initiator, by a Bus Device Reset message from any initiator, or a “hard” Reset condition. The drive shall ignore any attempt to release the reservation made by any other initiator.

5.2.1.8.4 Superseding Reservations

An initiator which holds a current reservation may modify that reservation by issuing another Reserve command to the same logical unit. The superseding Reserve command shall release the previous reservation state when the new reservation request is granted. The previous reservation shall not be modified if the new reservation request cannot be granted.

5.2.1.9 Release Command (17h)

The Release command (Table 5.2.1-17) is used to release previously reserved drives. It is not an error for an initiator to attempt to release a reservation that is not currently active. In this case, the drive returns Good status without altering any other reservation.

This command is implemented by the drive for an Entire Unit Release with Third Party Release supported and with the drive specific parameters listed in Table 5.2.1-17.

Table 5.2.1-17. Release Command

Bit Byte(s)	7	6	5	4	3	2	1	0
0	0	0	0	1	0	1	1	1
1	Logical Unit No. 0 0 [1]		0	3rd Pty [2]	Third Party Dev. ID [2]			Extent [3]
2	Reservation Identification 0 0 0 0 0 0 0 0 [4]							
3	Reserved 0 0 0 0 0 0 0 0							
4	Reserved 0 0 0 0 0 0 0 0							
5	0	0	0	0	0	0	Flag	Link [5]

Notes.

[1] In the CDB the LUN must be zero.

[2] Same as Note 2 for Reserve Command, Table 5.2.1-16.

[3] The Extent bit must be zero, if not supported by the drive. See drive Product Manual, section "SCSI Interface commands supported".

[4] Must be zero if not supported by the drive. See drive Product Manual, section "SCSI Interface commands supported". (see also paragraphs 5.2.1.8.1, 5.2.1.8.2).

[5] See "Control Byte" paragraph 4.2.6.

5.2.1.9.1 Logical Unit Release

If the extent bit is zero, the Release command shall cause the drive to terminate all reservations from the initiator to the drive.

5.2.1.9.2 Extent Release

The drive may not support extent reservations, in which case this bit must always be zero. Since the Reservation Identification byte is valid only for extent reservations, the drive ignores this byte when it does not support extent release. Check the drive Product Manual section "SCSI Interface commands supported" for support of this command.

5.2.1.9.3 Third Party Release

The drive supports the third party release option. The third party release option for the Release command allows an initiator to release a logical unit which was previously reserved using the third party reservation option (see Section 5.2.1.8.3).

If the third party (3rdPty) bit is zero, the third party release option is not requested. If the 3rdPty bit is one, the drive shall release the specified logical unit, but only if the reservation was made using the third party reservation option by the initiator that is requesting the release, and for the same SCSI device specified in the third party ID field.

5.2.1.10 Mode Sense (6) Command (1Ah)

Table 5.2.1-18. Mode Sense Command

Bit Byte(s)	7	6	5	4	3	2	1	0
0	0	0	0	1	1	0	1	0
1	Logical Unit No. [1] 0 0 0			0	DBD 0 [2]	0	0	0
2	PCF (See Table 5.2.1-19)[3]		Page Code				[4]	
3	0	0	0	0	0	0	0	0
4	Allocation Length						[5]	
5	0	0	0	0	0	0	Flag	Link[6]

The Mode Sense command provides a means for the drive to report its medium, logical unit, or peripheral device parameters to the initiator. It is a command complementary to the Mode Select command.

Notes.

This command is implemented in the drive with the following drive specific parameters:

[1] The LUN must be zero.

[2] A DBD bit of zero indicates that the drive may return zero or more block descriptors in the returned Mode Sense data, at the drive's discretion. Seagate SCSI-2 products return one block descriptor if the DBD bit is zero. A DBD bit of one specifies that the drive shall not return any block descriptors in the returned Mode Sense data. Byte 3 of the Mode Sense header contains 00h to indicate a block descriptor length of zero.

[3] The content of mode parameter bytes is determined by the value of the PCF (Page Control Field) bits specified in CDB byte 2, bits 6 & 7. The drive shall return the same Page Length for each supported page regardless of the value of PCF. The PCF field is defined in Table 5.2.1-19.

Table 5.2.1-19. Page Control Field Bits Affect On Mode Parameters Returned

PCF Bit		
7	6	
0	0	Return Current values. The Current values are the values currently being used by the drive to control its operation. After a Power On Reset, a hard Reset, or a Bus Device Reset message the Current values are equal to the Saved values (if Saved values can be retrieved) or the Default values (if Saved values cannot be retrieved). The Current value of a parameter is updated by a Mode Select command if the Mode Select Command ends with Good status returned.
0	1	Return Changeable values. The changeable values of any page is a mask that indicates the parameters that shall be changed via a Mode Select command and the parameters that shall not. Each returned parameter byte shall contain ones where a field or bit may be changed and zeros where a field or bit may not be changed.
1	0	Return Default values. The Default values are the values to which the drive sets the Current values after a reset condition unless valid Saved values are available.
1	1	Return Saved values. The saved values are the values the drive stores in nonvolatile memory. The Saved values of any changeable parameter can be set to new values via a Mode Select command with the SMP bit set to one. For nonchangeable parameters, the Default value is used.

The Block descriptor contains its normal values regardless of the value of the PCF. Unsupported fields or bits within a page are returned as zeros for all PCF values.

- [4] The Page Code allows the initiator to select one or all of the pages of Mode parameters supported by the target. Page Codes that may be supported by the drive are summarized here (see individual drive Product Manual):

Page Code	Description
0h	Unit Attention Page parameters (returned last of the pages).
01h	Error Recovery parameters.
02h	Disconnect/Reconnect Control parameters.
03h	Format parameters.
04h	Rigid Drive Geometry parameters.
07h	Verify Error Recovery Page parameters.
08h	Caching page
0Ah or 1Ah	Control Mode page
0Ch	Notch and Partition page
0Dh	Power condition page
10h	Xor Control mode page
1Ch	Information Exceptions Control page
3Fh	Return all supported pages.

- [5] The Allocation Length specifies the number of bytes that the initiator has allocated for returned Mode Sense data. An Allocation Length of zero indicates that no Mode Sense data shall be transferred. This condition shall not be considered as an error. Any other value indicates the *maximum* number of bytes that shall be transferred. The drive shall terminate the Data In phase when allocation length bytes have been transferred or when all available Mode Sense data has been transferred to the initiator, whichever is less.

- [6] See "Control Byte" paragraph 4.2.6.

Table 5.2.1-20. Mode Sense Data

Bit Byte(s)	7	6	5	4	3	2	1	0
0	Sense Data Length							[1]
1	0	0	0	0	0	0	0	0 [2]
2	WP [3]	Reserved		DPO-FUA [8]	Reserved			0
3	0	0	0	0	1	0	0	0 [4]

BLOCK DESCRIPTOR DATA

0	Density Code (MSB)							[5][6]
1	Number of Blocks (MSB)							[6]
2	Number of Blocks							[6]
3	Number of Blocks (LSB)							[6]
4	0	0	0	0	0	0	0	0
5	Block Length (MSB)							[7]
6	Block Length							[7]
7	Block Length (LSB)							[7]

PARAMETER INFORMATION

0-n	Mode Sense Page Headers and Their Parameters
-----	--

See following page for notes.

continued from previous page

- [1] The Sense Data Length specifies the length in bytes of the following Mode Sense data that is available to be transferred during the Data In phase. The Sense Data Length does not include itself.
- [2] The drive supports only 00h (default medium) in the Medium Type field.
- [3] A Write Protect (WP) bit of zero indicates the medium is write enabled. A WP bit of one indicates the medium is write protected.
- [4] The Block Descriptor Length specifies the length in bytes of the Block Descriptor. It is equal to the number of bytes in the Block Descriptor (8) and does not include the page headers and mode parameters, if any. The drive sends one Block Descriptor.

Each Block Descriptor specifies the medium characteristics for all or part of a logical unit. Each Block Descriptor contains a Density Code, a Number of Blocks, and a Block Length.

- [5]
 - a. Drives that do not support capacity programming have only 00h (default density) in the Density Code field.
 - b. Drives that support capacity programming may have a value in this field that states either the logical or actual capacity of the drive.
- [6]
 - a. For drives that do not support capacity programming, the Number of Blocks field specifies the number of logical blocks of the medium that meets the Density Code and Block Length in the Block Descriptor. A Number of Blocks of zero indicates that all of the remaining logical blocks of the logical unit have the medium characteristics specified by the Block Descriptor.
 - b. For drives that have capacity programming capability, these bytes do not report back the drive capacity on some models. These bytes are always zero on those models. Some models do report drive capacity in bytes 0, 1, 2 and 3, so those bytes will be non-zero. See the Product Manual Volume I for the drive of interest.
- [7] The Block Length, as defined after a format function, specifies the length in bytes of each logical block described by the Block Descriptor. Default is 512 if no Mode Select command is received before the Format command. The usual valid values are 256 through 4096. Some drive products can format down to 180 bytes per sector. Some products can only format an even numbered value of bytes per sector (180 - 4096).
- [8] When used with the Mode Sense command, a DPOFUA bit of zero indicates that the target does not contain a cache memory or does not support the DPO and FUA bits. A DPOFUA bit of one indicates that the target supports the DPO and FUA bits.

Mode Sense Page Descriptor Header

Table 5.2.1-21. Mode Sense Page Descriptor Header

Bit Byte(s)	7	6	5	4	3	2	1	0
0	PS	0	Page Code					[1]
1	Page Length							[1]
2-n	Mode Parameters							[1]

[1] Page Code (Hex)	Page Description	Reference Table	Page
01	Error Recovery	5.2.1-22	155
02	Disconnect/Reconnect Control	5.2.1-23	158
03	Format Parameters	5.2.1-24	161
04	Rigid Drive Geometry	5.2.1-25	164
07	Verify Error Recovery Page	5.2.1-26	167
08	Caching Page	5.2.1-27	169
0A	Control Mode Page	5.2.1-28/29	172/173
0C	Notch and Partition Page	5.2.1-30	174
0D or 1A	Power Condition Page	5.2.1-31	176
10	Xor Control Mode page	5.2.1-32	177
00	Unit Attention Page Parameters	5.2.1-33	178
1C	Information Exceptions Control page	5.2.1-33a	180

Each page of mode parameters (for the Mode Sense command) begins with a two byte Page Descriptor Header. The Page Code identifies the page of mode parameters that is being transferred. The Page Length indicates the number of additional bytes of mode parameters being sent by the drive. The parameter bit values are left blank herein, because they may be different for each drive model.

Note: See the individual drive Product Manuals (Volume 1) for a table giving the Mode Sense parameter values that are applicable to the drive model of interest. The tables in the Volume 1 also show which parameters are changeable in the drive model of interest and which are not.

Multiple pages of mode parameters may be transferred in one Mode Sense Data In phase (using Page Code 3Fh). If a non-supported page code is requested by the Initiator, the drive terminates the command with CHECK CONDITION status, sets the sense key to 05, ILLEGAL REQUEST, and sets the additional sense code to 24, INVALID FIELD IN PARAMETER LIST.

The Parameters Savable (PS) bit, when set to one, indicates if the page contains savable parameters. When the PS bit is set to zero, none of the parameters within the page are savable. Since the parameters within pages 3 and 4 are always saved during Format commands (but not via a Mode Select command with the SMP bit set to 1), these pages return a one for the PS bit.

Error Recovery Page (01)

The drive Error Recovery Page implementation is defined in Table 5.2.1-22. This table summarizes: the function, the default value, and changeability status for each byte/bit.

Table 5.2.1-22. Error Recovery Page**PAGE DESCRIPTOR HEADER**

Bit Byte(s)	7	6	5	4	3	2	1	0
0	PS 1 [1]	0	0	0	0	0	0	1
1	Page Length (0Ah)							

ERROR RECOVERY PARAMETERS

2 DEFAULT	AWRE	ARRE	TB	RC	EER	PER	DTE	DCR [3]
CHANGEABLE	[2]							
3 DEFAULT	Read Retry Count [4]							
CHANGEABLE	[2]							
4 DEFAULT	Correction Span (bits) [5]							
CHANGEABLE	[2]							
5 DEFAULT	Head Offset Count [6]							
CHANGEABLE	[2]							
6 DEFAULT	Data Strobe Offset Count [7]							
CHANGEABLE	[2]							

ERROR RECOVERY PARAMETERS (continued)

Bit Byte(s)	7	6	5	4	3	2	1	0
7 DEFAULT	Reserved							[10]
CHANGEABLE	[2]							
8 DEFAULT	Write Retry Count							[8]
CHANGEABLE	[2]							
9 DEFAULT	Reserved							
CHANGEABLE	[2]							
10	(MSB)	Recovery Time Limit						[9]
11								LSB
CHANGEABLE	[2]							

- [1] The returned PS (Parameter Savable) bit of 1 indicates that page 01h parameter data is savable. This bit is not used with the Mode Select command.
- [2] A value of zero means this bit function is not directly changeable by an initiator, a value of 1 means the bit function is directly changeable by an initiator. (See Mode Select Command). See drive Product Manuals (Volume 1) section showing changeable values.
- [3] The Automatic Write Reallocation of defective data blocks Enabled (AWRE) bit, when set to one, allows the drive to automatically relocate bad blocks detected during write operations. The drive performs the automatic write reallocation only if the drive has the valid data (e.g., original data in the buffer or recovered from the medium). The valid data is placed in the reallocated block. This function doesn't apply to the Format Unit command. When set to zero, the drive shall not perform automatic reallocation but shall create Check Condition status with Sense Key of Medium Error instead.

The Automatic Read Reallocation of defective data blocks Enabled (ARRE) bit, when set to one, allows the drive to automatically relocate bad blocks detected during read operations. Automatic reallocation is performed only if the drive successfully recovers the data and is able to place it in the reallocated block. When set to zero, the drive shall not perform automatic reallocation but shall create Check Condition status with Sense Key of Medium Error instead.

The Transfer Block (TB) bit, when set to one, indicates the data block that is not recovered shall be transferred to the initiator. When set to zero, the failing data block shall not be transferred.

The Read Continuous (RC) bit, when set to one, requests the drive to transfer the requested data length without adding delays (for retries or ECC correction) that may be required to ensure data integrity. The drive may send erroneous data in order to maintain the continuous flow of data. This bit shall override the DTE bit if it is set. RC bit has priority also over ARRE, AWRE, EER, DCR and PER bits. If the RC bit is set to one, no Auto Reallocation will be attempted. When set to zero, recovery actions during data transfer are allowed. This bit is set to zero and is not changeable in most if not all of the drive models covered by this manual. See drive Product Manual Volume 1 section showing changeable values.

continued from previous page

The Enable Early Recovery (EER) bit, when set to one, allows the drive to apply on-the-fly T>1 ECC correction as soon as possible, before attempting other retry mechanisms, and without reporting successful corrections to the host as recovered errors. Seek error retries and message system errors are not affected by this bit. When this bit is set, the DCR bit must be zero. When the EER bit is set to zero, the drive shall apply ECC correction before other retry mechanisms, but shall not perform T>1 corrections on the fly; any successful correction will be reported to the host as a recovered error.

The Post Error (PER) bit, when set to one, indicates the drive reports Check Condition status and appropriate Sense Key for any *recovered errors* encountered. Reporting of unrecoverable errors has priority over reporting of recoverable errors. When set to zero, any errors recovered within the limits established by the other Error Recovery Flags are not reported. Any unrecoverable errors are reported.

The Disable Transfer on Error (DTE) bit is valid only when the PER bit is set to one. When the DTE bit is set to one, it indicates the drive terminates data transfer even for recoverable errors (the drive will transfer the data for the recovered error before terminating the transfer). When DTE is set to zero, data transfer continues if recoverable errors are encountered. If the PER bit is one and the DTE bit is zero, recoverable errors are reported after all data has been transferred.

The Disable Correction (DCR) bit, when set to one, indicates ECC correction shall not be applied to the data even if correction is possible. When set to zero, ECC correction shall be applied if correction is possible.

- [4] The Read Retry Count sets up the maximum amount of error recovery effort to be applied for each LBA that could not be recovered during a read operation. The hex value in this field specifies the maximum error recovery level that the drive applies during a read operation to the recovery of an LBA needing recovery effort. Each level may consist of multiple error recovery steps. See the individual Product Manuals for more details on the levels of error recovery available.
- [5] The Correction Span is the size of the largest read data error, in bits, on which ECC correction is to be attempted. Errors longer than this span are reported as unrecoverable. This value is drive dependent. If this field is zero the drive uses its default value.
- [6] The drive Head Offset Count is a default of zero and not changeable to signify that this feature is not programmable by the initiator. Head offsets are performed as part of the drive's retry algorithms.
- [7] The drive Data Strobe Offset Count is a default of zero and not changeable to signify that this feature is not programmable by the initiator. Data Strobe Offsets are performed as part of the drive's retry algorithms.
- [8] The Write Retry Count sets up the maximum amount of error recovery to be applied for each LBA that could not be recovered during a write operation. The hex value in this field specifies the maximum error recovery level that the drive applies during a write operation to the recovery of an LBA needing recovery effort. Each level may consist of multiple error recovery steps. See the individual Product Manuals for more details on the levels of error recovery available.
- [9] The Recovery Time Limit field (bytes 10 and 11) specifies the maximum time in milliseconds that the host allows the drive to spend in error recovery efforts during the execution of a command. The Read and Write Retry count can also be set to limit the amount of time the drive spends in error recovery of individual LBAs. The total of all times used to recover individual LBAs in the block called for by a command cannot exceed the Recovery Time Limit value in bytes 10 and 11. Once the drive has reached the error recovery time limit for a particular command, the command ends with a CHECK CONDITION status and an unrecovered error is reported. A Recovery Time Limit of FFFFH means that the command recovery time is unlimited. A value of 0000H means that no time shall be spent in error recovery. A changeable Recovery Time Limit is not supported on all drives supported by this manual. See individual Product Manual Mode page changeable bit settings for Mode page 01h, bytes 10 and 11.
- [10] In SCSI-1 mode of operation, this byte is the recovery time limit value (see [9] above).

Disconnect/Reconnect Control Page (02)

The Disconnect/Reconnect Page implementation is defined in Table 5.2.1-23. This table summarizes the function and defines the default value and changeable status.

Table 5.2.1-23. Disconnect/Reconnect Control Page

PAGE DESCRIPTOR HEADER								
Bit Byte(s)	7	6	5	4	3	2	1	0
0	PS		Page Code 02h					
1	1[1]	0	0	0	0	0	1	0
			Page Length (0Eh)					
DISCONNECT/RECONNECT CONTROL PARAMETERS								
2	Buffer Full Ratio							[3]
DEFAULT								
CHANGEABLE	[2]							
3	Buffer Empty Ratio							[4]
DEFAULT								
CHANGEABLE	[2]							
4	Bus Inactivity Limit (MSB)							[5]
DEFAULT								
CHANGEABLE	[2]							
5	Bus Inactivity Limit (LSB)							
DEFAULT								
CHANGEABLE	[2]							
6,7	Disconnect Time Limit							[6]
DEFAULT								
CHANGEABLE	[2]							
8,9	Connect Time Limit							[7]
DEFAULT								
CHANGEABLE	[2]							
10,11	MSB	Maximum Burst Size					[8]	
DEFAULT							LSB	
CHANGEABLE	[2]							
12	EMDP	0	0	0	DImm	DTDC[9]		
DEFAULT	[10]				[11]			
CHANGEABLE	[2]							
13 - 15	Reserved							

Notes for Table 5.2.1-23.

- [1] The PS (Parameter Savable) bit of 1 indicates that the page 02h parameter data is savable.
- [2] A changeable value of zero means this function is not directly changeable by an initiator. A value of 1 means the bit function is directly changeable by an initiator. (See Mode Select command.) See drive Product Manual Mode Sense data section for table showing changeable values.
- [3] Both ratio parameters are the numerators of a fractional multiplier that has 256 (100h) as its denominator.
- [4] inator.

continued from previous page

- [3] The Buffer Full Ratio indicates, on Read commands, how full the drive's buffer shall be before attempting a reselection. The drive rounds the requested ratio up to the nearest whole logical block. This value is changeable by an initiator.

The buffer full ratio is taken to be a percentage of the smaller of

- a) the buffer size
- or
- b) the remaining transfer length.

For example, if the buffer full ratio is 80h (128 Decimal) (indicating a 128/256 or 50% value), the transfer length of a read command is 20h blocks, and the buffer size is 30h blocks, the reconnect begins when 10h blocks (50% of the transfer length of 20h blocks) is in the buffer.

- [4] The Buffer Empty Ratio indicates, on Write commands, how empty the drive's buffer shall be before attempting a reselection. The drive, rounds the requested ratio down to the nearest whole logical block. This value is changeable by an initiator.
- [5] The Bus Inactivity Limit field (bytes 4 & 5) indicates the time, in 100 microsecond increments, the drive is allowed to assert the Busy signal without handshakes until it shall disconnect. The value of ten indicates the drive is allowed to maintain the Busy signal for one millisecond without handshakes. A value of zero indicates that there is no bus inactivity limit. This value is not changeable by the initiator.
- [6] The Disconnect Time Limit field (bytes 6 and 7) indicates the minimum time, in 100 microsecond increments, the drive shall remain disconnected until it shall attempt to reconnect. A value of zero indicates the drive is allowed to reconnect immediately. For the typical drive, this is likely to always be zero and the changeable code is always zero. (See individual Product Manual).
- [7] The Connect Time Limit field (bytes 8 and 9) indicates the maximum time in 100 microsecond increments that the target should remain connected until it attempts to disconnect. A setting of zero indicates that the drive is allowed to remain connected indefinitely until it attempts disconnection.
- [8] The maximum burst size field indicates the maximum amount of data that the drive shall transfer during a data phase before disconnecting if the initiator has granted the disconnect privilege. This value is expressed in increments of 512 bytes (e.g., a value of one means 512 bytes, two means 1024 bytes, etc.). A value of zero indicates there is no limit on the amount of data transferred per connection.

continued from previous page

- [9] The data transfer disconnect control (DTDC) field is intended to define further restrictions on when a disconnect is permitted. The various DTDC functions called for by the DTDC Field Codes are given in the table following.

Data Transfer Disconnect Control

DTDC	Description
000b	Data transfer disconnect control is not used. Disconnect is controlled by the other fields in this page.
001b	A drive shall not attempt to disconnect once the data transfer of a command has started until all data the command is to transfer has been transferred. The connect time limit and bus inactivity limit are ignored during the data transfer.
010b	Reserved
011b	A drive 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 data transfer has started.

If DTDC is nonzero and the maximum burst size is nonzero the drive shall return CHECK CONDITION status. The sense key shall be set to ILLEGAL REQUEST and the additional sense code set to ILLEGAL FIELD IN PARAMETER LIST.

Not all Seagate drive models implement this field. See drive's Product Manual Volume 1, which indicates if a particular drive implements the DTDC function.

- [10] The Enable Modify Data Pointers (EMDP) bit indicates whether or not the initiator allows the Modify Data Pointers message to be sent by the target. If the EMDP bit is zero, the target shall not issue the Modify Data Pointers Message. If the EMDP bit is one, the target is allowed to issue Modify Data Pointers Message.

- [11] DImm: Disconnect Immediate.

The optional Disconnect Immediate (DImm) bit of zero indicates that the drive may disconnect after command phase if it chooses to do so based on its internal algorithms, the setting of the DiscPriv bit in the Identify message and the settings of the other parameters in this mode page.

An optional Disconnect Immediate bit of one indicates that the drive shall attempt to disconnect immediately after every command phase for those connections in which disconnections are allowed.

Host adapters in untagged command environments which have relatively high overhead to handle a Disconnect message sequence may want to use the Disconnect Immediate feature. However, not all Seagate drive models implement this feature. See individual Product Manual, Volume 1.

Format Parameters Page (03)

The Format Parameters Page implementation is defined in Table 5.2.1-24. This table summarizes the function and defines the default or changeability status for each bit.

Table 5.2.1-24. Format Parameter Page

Page Descriptor Header								
Bit Byte(s)	7	6	5	4	3	2	1	0
0	PS 1[1]	0	Page Code 03h 0 0 0 0 1 1					
1	Page Length (16h)							
Format Parameters [1]								
2,3 DEFAULT	Tracks per Zone (MSB)							
CHANGEABLE	[12]							
4,5 DEFAULT	Alternate Sectors per Zone							
CHANGEABLE	[12]							
6,7 DEFAULT	Alternate Tracks per Zone							
CHANGEABLE	[12]							
8,9 DEFAULT	Alternate Tracks per Volume							
CHANGEABLE	[12]							
10,11 DEFAULT	Sectors per Track							
CHANGEABLE	[12]							
12,13 DEFAULT	Data Bytes per Physical Sector							
CHANGEABLE	[12]							
14,15 DEFAULT	Interleave							
CHANGEABLE	[12]							
16,17 DEFAULT	Track Skew Factor							
CHANGEABLE	[12]							
18,19 DEFAULT	Cylinder Skew Factor [10]							
CHANGEABLE	[12]							
20 DEFAULT	SSEC	HSEC	RMB	SURF	Reserved			[11]
CHANGEABLE	[12]							
21,22,23 DEFAULT	Reserved							
CHANGEABLE	[12]							

Notes for Table 5.2.1-24. Format Parameter Page (continued)

- [1] The only time this page of parameters may be sent is immediately before sending a Format Unit command to the drive. The Current parameters for this page are updated immediately but any changes between these Current parameters and the existing media format are not in effect until after the Format Unit command is completed. A PS bit of 1 indicates this page is savable. The PS bit is not used with the Mode Select command.
- [2] The Tracks per Zone* field indicates the number of tracks the drive allocates to each defect management zone. A zone can be one or more tracks or one or more cylinders. See individual drive Product Manual for number of tracks allocated to each defect management zone for that drive model.
- [3] The Alternate Sectors per Zone* field indicates the number of spare sectors to be reserved for the defined defect management zone. A value of zero indicates that no sectors are to be reserved in each zone for defect management. This is to accommodate hosts that want to manage the defects themselves.
- [4] The Alternate Tracks per Zone* field indicates the number of spare tracks to be reserved at the end of each defect management zone. A value of zero indicates that no spare tracks are to be reserved in each zone for defect management by the drive.
- [5] The Alternate Tracks per Volume field indicates the number of spare tracks to be reserved at the end of the drive volume. The drive uses these locations for replacing defective sectors. A value of zero indicates that no spare tracks are to be reserved at the end of the unit for defect management. The initiator may change this value for a number between 0 and 255 that is a multiple of the total number of Data Read/Write heads installed. However, it is not changeable on some products.
- [6] The Sectors per Track field indicates the *average* number of physical sectors the drive has per disc track. This value depends on the selected sector size and ZBR zones. The number of user accessible sectors per track may be fewer than the reported value, since sectors per Track includes sectors set aside for defect management. This value cannot be used to calculate drive user accessible capacity.

Note. The value cannot be directly selected with the Mode Select command, but is a report of how the drive is configured.
- [7] The Bytes per Physical Sector field indicates the number of data bytes the drive shall allocate per physical sector. This value equals the block length reported in the Mode Sense block descriptor. The bytes per physical sector is not directly changeable by the initiator and is not verified on a Mode Select command.

* Defect management zone (one or more tracks), NOT a ZBR (variable track capacity recording) zone. ZBR zones are referred to as notches (page 0Ch is the Notch page).

The actual implementation of reserving spare areas for defect management takes place during the Format Unit command.

continued from previous page

[8] The Interleave field is the interleave value sent to the drive during the last Format Unit command.

Note. This field is valid only for Mode Sense commands. The drive ignores this field during Mode Select commands.

[9] The Track Skew Factor field indicates the average number of physical sectors between the last logical block on one track and the first logical block on the next sequential track of the same cylinder. A value of zero indicates no skew.

Note. This value is not changeable by an initiator.

[10] The Cylinder Skew Factor field indicates the average number of physical sectors between the last logical block of one cylinder and the first logical block of the next cylinder. A value of zero indicates no skew. Cylinder skew will be utilized by a drive but is not changeable by an initiator.

[11] The Drive Type field bits are defined as follows:

The Hard Sectoring (HSEC) bit (bit 6) set to one indicates the drive shall use hard sector formatting. The soft sectoring (SSEC) bit (bit 7) when set to one indicates the drive uses soft sectoring.

Bits 0-5, and 7 are not implemented by the drive and are always zero. All bits (0-7) are not changeable.

[12] See drive Product Manual Mode Sense Data section for changeable values.

Rigid Drive Geometry Parameters (04)

The Rigid Drive Geometry Parameters Page implementation is defined in Table 5.2.1-25. This table summarizes the function and defines the default value.

Table 5.2.1-25. Rigid Drive Geometry Parameters**PAGE DESCRIPTOR HEADER**

Bit Byte(s)	7	6	5	4	3	2	1	0
0	PS 1[1]	0	0	0	0	1	0	0
1	Page Length (16h)							

RIGID DRIVE GEOMETRY PARAMETERS

2 DEFAULT CHANGEABLE	Number of Cylinders (MSB) [2]							
3 DEFAULT CHANGEABLE	Number of Cylinders [8]							
4 DEFAULT CHANGEABLE	Number of Cylinders (LSB) [2]							
5 DEFAULT CHANGEABLE	Number of Heads [3]							
6,7,8 DEFAULT CHANGEABLE	Starting Cylinder - Write Precomp [4]							
9,10,11 DEFAULT CHANGEABLE	Starting Cylinder-Reduced Write Current [4]							
12,13 DEFAULT CHANGEABLE	Drive Step Rate [4]							
14,15,16 DEFAULT CHANGEABLE	Landing Zone Cylinder [4]							

PAGE DESCRIPTOR HEADER (continued)

Bit Byte(s)	7	6	5	4	3	2	1	0
17 DEFAULT	0	0	0	0	0	0	RPL	[5]
CHANGEABLE	-----[8]-----							
18 DEFAULT	Rotational Offset							
CHANGEABLE	XXh -----[6]-----							
19 DEFAULT	Reserved							
CHANGEABLE	0	0	0	0	0	0	0	0
20 21	(MSB) ----- Medium Rotation Rate ----- [7] ----- (LSB)							
CHANGEABLE	-----[8]-----							
22 DEFAULT	Reserved							
CHANGEABLE	0	0	0	0	0	0	0	0
23 DEFAULT	Reserved							
CHANGEABLE	0	0	0	0	0	0	0	0
	-----[8]-----							

Notes for Table 5.2.1-25.

- [1] The PS bit of 1 indicates that page 04 parameter data is savable and is saved when a Format Function is performed. In some drives an exception exists that applies to bytes 17 and 18. In the exception drives, bytes 17 and 18 are only saved if the SMP bit in the Mode Select Command (Table 5.2.1-12) is 1. See applicable drive Product Manual Volume 1, "SCSI Bus conditions and miscellaneous features supported". This PS bit is not applicable to the Mode Select command.
- [2] The number of cylinders field defines the number of physical cylinders used for data storage. This may or may not include spare cylinders set aside for flaw reallocation. See individual drive Product Manual Volume 1 which specifies what the drive reports.

The drive uses some additional cylinders for storing drive parameters, defect lists, or for diagnostic purposes. These are not accessible by the user.

- [3] The Number of Heads field indicates the maximum number of data (read/write) heads on the drive.
- [4] Not applicable.
- [5] Used for Spindle Synchronization (rotation position locking).

Rotational Position Locking

RPL	Description
00b	Indicates that spindle synchronization is automatic. (Automatic master arbitration is used to determine which device in the chain is to be master).
01b	The target operates as a synchronized-spindle slave.
10b	The target operates as a synchronized-spindle master.
11b	The target operates as a synchronized-spindle master control (Not supported by drive).

- [6] Rotational skew in the lagging direction used for spindle synchronization. The value XXh given represents a XXh/FFh fractional part of a revolution lagging offset. One revolution lag is maximum. See section on Synchronous Spindle Operation in drive Product Manual Vol. 1.
- [7] On Mode Sense Command these bytes return drive nominal rotation rate in revolutions per minute for synchronous spindle operation. The bytes have no meaning for Mode Select.
- [8] See drive Product Manual Mode Sense Data section for changeable values.

Table 5.2.1-26. Verify Error Recovery Page (07h)
PAGE DESCRIPTOR HEADER

Bit Byte(s)	7	6	5	4	3	2	1	0
0	PS 1[1]	Rsvd 0	Page Code (07h)					
1	Page Length (0Ah)							

VERIFY ERROR RECOVERY PARAMETERS

	7	6	5	4	3	2	1	0	
2	Reserved				ERR	PER	DTE	DCR[2]	
DEFAULT	0	0	0	0					
CHANGEABLE	[6]								
3	Verify Retry Count [3]								
DEFAULT									
CHANGEABLE	[6]								
4	Verify Correction Span (Bits) [4]								
DEFAULT									
CHANGEABLE	[6]								
5	Reserved								
DEFAULT	0	0	0	0	0	0	0	0	
CHANGEABLE	[6]								
6	Reserved								
DEFAULT	0	0	0	0	0	0	0	0	
CHANGEABLE	[6]								
7	Reserved								
DEFAULT	0	0	0	0	0	0	0	0	
CHANGEABLE	[6]								
8	Reserved								
DEFAULT	0	0	0	0	0	0	0	0	
CHANGEABLE	[6]								
9	Reserved								
DEFAULT	0	0	0	0	0	0	0	0	
CHANGEABLE	[6]								
10	(MSB)	Verify Recovery Time Limit							[5]
DEFAULT	1	1	1	1	1	1	1	1	
CHANGEABLE	[6]								
11	Verify Recovery Time Limit							(LSB)	
DEFAULT	1	1	1	1	1	1	1	1	
CHANGEABLE	[6]								

The verify error recovery page (Table 5.2.1-26) specifies the error recovery parameters the target shall use during the Verify command and the verify operation of the Write and Verify command.

Notes.

- [1] The parameters savable (PS) bit is used only with the Mode Sense command. This bit is reserved with the Mode Select command. A PS bit of one indicates that the target is capable of saving the page in a non-volatile vendor specific location.
- [2] The EER, PER DTE and DCR bits are defined in Table 5.2.1-22 notes.
- [3] The Verify Retry Count sets up the maximum amount of error recovery effort to be applied for each LBA that could not be recovered during a verify operation. The hex value in this field specifies the maximum error recovery level that the drive applies during a verify operation to the recovery of an LBA needing recovery effort. Each level may consist of multiple error recovery steps. See the individual Product Manuals for more details on the levels of error recovery available.
- [4] The verify correction span field specifies the size, in bits, of the largest burst data error for which data error correction may be attempted. If the drive does not implement this field, a value of zero is returned in Mode Sense data.
- [5] The Recovery Time Limit field (bytes 10 and 11) specifies the maximum time in milliseconds that the host allows the drive to spend in error recovery efforts during the execution of a command. The Verify Retry count can also be set to limit the amount of time the drive spends in error recovery of individual LBAs. The total of all times used to recover individual LBAs in the block called for by a command cannot exceed the Recovery Time Limit value in bytes 10 and 11. Once the drive has reached the error recovery time limit for a particular command, the command ends with a CHECK CONDITION status and an unrecovered error is reported. A Recovery Time Limit of FFFFH means that the command recovery time is unlimited. A value of 0000H means that no time shall be spent in error recovery. A changeable Recovery Time Limit is not supported on all drives supported by this manual. See individual Product Manual Mode page changeable bit settings for Mode page 01h, bytes 10 and 11.

To disable all types of correction and retries the initiator sets the ERR bit to zero, the PER, DTE and DCR bits to one and the number of retries and recovery time limit to zero.
- [6] See drive Product Manual Mode Sense Data section for changeable values.

Caching Page For Mode Sense/Mode Select (08)

The caching parameters page (08h) defines the parameters that affect the use of the cache. See Table 5.2.1-27.

Table 5.2.1-27. Caching Parameters**PAGE DESCRIPTOR HEADER**

Bit Byte(s)	7	6	5	4	3	2	1	0
0	PS[1]	Rsvd	Page Code (08h)					
1	Page Length (12h)							
2	IC	ABPF	CAP	DISC	SIZE	WCE	MF	RCD [3]
CHANGEABLE	[2]							
3	Demand Read Retention Pri. [4]				Write Retention Priority [5]			
CHANGEABLE	[2]				[2]			
4	(MSB) Disable Pre-Fetch Transfer Length [6]							
5	LSB							
CHANGEABLE	[2]							
6	(MSB) Minimum Pre-Fetch [7]							
7	LSB							
CHANGEABLE	[2]							
8	(MSB) Maximum Pre-fetch [8]							
9	LSB							
CHANGEABLE	[2]							
10	(MSB) Maximum Pre-fetch Ceiling [9]							
11	LSB							
CHANGEABLE	[2]							
12	FSW[10]	LBCSS[15]	DRA[11]	RSVD				
CHANGEABLE	[2]							
13	Number of Cache Segments [12]							
CHANGEABLE	[2]							
14	(MSB) Cache Segment Size [13]							
15	LSB							
CHANGEABLE	[2]							
16	Rsvd							
17	(MSB) Non-Cache Segment Size [14]							
18	LSB							
19	LSB							
CHANGEABLE	[2]							

Notes for Table 5.2.1-27.

- [1] The returned PS (Parameter Savable) bit of 1 indicates that page 01h parameter data is savable.
- [2] A value of zero means this bit function is not directly changeable by an initiator, a value of 1 means the bit function is directly changeable by an initiator. (See Mode Select Command). See drive Product Manuals for changeable values.
- [3] IC: The Initiator Control (IC) enable bit (Bit 7 Byte 2), when set to one, adaptive read look-ahead (ARLA) is disabled. When IC is set to zero (ARLA) is enabled. Since Seagate drives covered by this manual never organize the cache according to size of segment, but rather by number of segments, this bit is used to enable or disable adaptive RLA (in non-Seagate equipment, this might be used to designate cache size).

Note. ARLA cannot be disabled in some Seagate drives using the ASAll code. See individual drive Product Manual (Vol. 1).

ABPF: The Abort Pre-Fetch (ABPF) bit (Bit 6 Byte 2), when set to one, with the DRA bit equal to zero, requests that the SCSI device abort the pre-fetch upon selection. The ABPF set to one takes precedence over the Minimum Pre-fetch bytes. When set to zero, with the DRA bit equal to zero, the termination of any active pre-fetch is dependent upon Caching Page bytes 4 through 11 and is operation and/or vendor specific.

CAP: Caching Analysis Permitted (CAP) a one in this bit enables caching analysis. A zero indicates caching analysis is disabled. Caching analysis results are placed in the SCSI Logging Information Table, Table 5.1.3.2-9. See individual drive Product Manual "SCSI Bus Conditions and Miscellaneous Features Supported" table.

DISC: The Discontinuity (DISC) bit (Bit 4 Byte 2), when set to one, requests that the SCSI device continue the pre-fetch across time discontinuities, such as across cylinders or tracks up to the limits of the buffer, or segment, space available for pre-fetch. When set to zero, the DISC requests that pre-fetched be truncated at time discontinuities.

SIZE: The Size Enable (SIZE) bit (Bit 3 Byte 2), when set to one, indicates that the Cache Segment Size is to be used to control caching segmentation. When SIZE equals zero, the Initiator requests that the *Number* of Cache Segments is to be used to control caching segmentation. For Seagate drives covered by this manual SIZE is always zero.

WCE: Write Cache Enable.

- 0: SCSI Write commands may not return status and completion message bytes until all data has been written to the media.
- 1: SCSI Write commands may return status and completion message bytes as soon as all data has been received from the host.

MF: Multiplication Factor.

- 0: The "Minimum Pre-fetch" and "Maximum pre-fetch" fields are interpreted as a number of logical blocks.
- 1: Specifies that the target shall interpret the minimum and maximum pre-fetch fields to be specified in terms of a scalar number which, when multiplied by the number of logical blocks to be transferred for the current command, yields the number of logical blocks for each of the respective types of pre-fetch.

RCD: Read Cache Disable.

- 0: SCSI Read commands may access the cache or the media.
- 1: SCSI Read commands must access the media. Data cannot come from the cache.

- [4] Demand Read Retention Priority:
The cache replacement algorithm does not distinguish between retention in the cache of host-requested data and prefetch data. Therefore, this half byte is always 0.

continued from previous page

- [5] Write Retention Priority:
The cache replacement algorithm does distinguish between retention in the cache of host-requested data and prefetch data. Therefore, this half byte is always 0.
- [6] Disable Pre-fetch Transfer Length:
Prefetch is disabled for any SCSI Read command whose requested transfer length exceeds this value.
- [7] Minimum Pre-fetch:
Specifies the minimum number sectors to pre-fetch, regardless of the delay it may cause to other commands.
- [8] Maximum Pre-fetch:
Specifies the maximum number of logical blocks that may be pre-fetched. The pre-fetch operation may be aborted before the maximum pre-fetch value is reached, but only if the minimum pre-fetch value has been satisfied.
- [9] The maximum Pre-fetch Ceiling field specifies an upper limit on the number of logical blocks computed as the maximum pre-fetch. If the Maximum Prefetch value is greater than the Maximum Pre-fetch ceiling the value is Truncated to the Maximum Pre-fetch Ceiling value.
- [10] FSW: The Force Sequential Write (FSW) bit (Bit 7 Byte 12). When set to one, indicates that multiple block writes are to be transferred over the SCSI bus and written to the media in an ascending, sequential, logical block order. When the FSW bit equals zero, the target is allowed to reorder the sequence of writing addressed logical blocks in order to achieve a faster command completion.
- [11] DRA: The Disable Read-Ahead (DRA) bit (Bit 5 Byte 12), when set to one, requests that the target not read into the buffer any logical blocks beyond the addressed logical block(s). When the DRA bit equals zero, the target may continue to read logical blocks into the buffer beyond the addressed logical block(s).
- [12] Number of Cache Segments:
The Number of Cache Segments byte (Byte 13) gives the number of segments into which the host requests the drive divide the cache.
- [13] Cache Segment Size:
The Cache Segment Size field (Bytes 14 and 15) indicates the requested segment size in Bytes. This standard assumes that the Cache Segment Size field is valid only when the SIZE bit is one.
- [14] Non-Cache Segment Size:
If the Non-Cache Buffer size field (Bytes 17-19) is greater than zero, this field specifies to the target the number of bytes the initiator requests that the target allocate for a buffer function when all other cache segments are occupied by data to be retained. If the number is at least one, caching functions in the other segments need not be impacted by cache misses to perform the SCSI buffer function. The impact of the Non-Cache Buffer Size equal 0 or the same of this field plus the Cache Segment Size greater than the buffer size is vendor specific.
- [15] LBCSS:
Logical block cache segment size. Not used at this time.

In addition to the caching control provided by the Caching Mode page, some 10-byte commands contain control bits DPO and FUA the intent of which is to override the cache page control bits.

continued from previous page

DPO: Disable Page Out.

If DPO = 1, the cache replacement algorithm should not replace existing cache data with the current data (if possible). If a cache segment must be overwritten, it should be made the LRU.

FUA: Force Unit Access.

Read: If FUA = 1, the requested data must be read from the media. If cache data which overlaps the request has not yet been written to the media, it should be written before the read is allowed to occur.

Write: If FUA = 1, all data must be written to the media before the SCSI operation returns the status and completion message bytes.

Mode Sense/Mode Select Control Mode page (0Ah)

Table 5.2.1-28. Control Mode Page

Bit Byte(s)	7	6	5	4	3	2	1	0
0	PS 1[1]	Rsvd	Page Code (0Ah)					
1	Page Length (0Ah)							
2	Reserved						GLTSD	RLEC
DEFAULT	0	0	0	0	0	0	[9]	[2]
CHANGEABLE	[7]							
3	Queue Algorithm Modifier				Reserved		QErr	DQue
DEFAULT	[3]				0 0		[4]	[5]
CHANGEABLE	[7]							
4	EECA	RAC	Reserved			Not Supported		
DEFAULT	[6]	[10]	0 0 0			0 0 0		
CHANGEABLE	[7]							
5	Reserved							
DEFAULT	0	0	0	0	0	0	0	0
CHANGEABLE	[7]							
6 & 7	Not Supported							
DEFAULT	0	0	0	0	0	0	0	0
CHANGEABLE	[7]							
8 & 9	Busy Timeout Period							
DEFAULT	[8]							
CHANGEABLE	[7]							
10	Reserved (All Zeros)							
11	Reserved (All Zeros)							

continued from previous page

The control mode page (Table 5.2.1-28) provides controls over several SCSI-2 features which are applicable to all device types such as tagged queuing, extended contingent allegiance, asynchronous event notification, and error logging.

Notes.

- [1] The PS (Parameter Savable) bit of 1 indicates that the page 0Ah parameter data is savable in non-volatile memory.
- [2] A report log exception condition (RLEC) bit of one specifies that the target shall report log exception conditions. A RLEC bit of zero specifies that the drive does not report Log exception conditions.
- [3] The queue algorithm modifier field specifies restrictions on the algorithm used for re-ordering commands that are tagged with the SIMPLE QUEUE TAG message.

Table 5.2.1-29. Queue Algorithm Modifier

Value	Definition	Value	Definition
0h	Guaranteed data integrity	2h - 7h	Reserved
1h	Unrestricted re-ordering allowed	8h - Fh	Vendor Specific

A value of 0h in the Queue Algorithm Modifier bits requires the device to order the actual execution sequence of the queued command such that data integrity is guaranteed at any time. This requires that, if the data transmission of a command was halted at any time, the final value of all data must have exactly the same value it would have had if the command had been executed without queuing. The guaranteed data integrity value (0h) of the queue algorithm modifier bits is the usual default value.

A value of 1h in the Queue Algorithm Modifier bits allows the device to order the actual execution sequence of the queued commands in any manner it selects. Any data integrity problems related to command sequence ordering are explicitly handled by the host operating system software.

- [4] A Queue Error Management (QErr) bit of zero indicates that those commands still queued after the device has entered the contingent allegiance condition continue execution in a normal manner when that condition has terminated. A QErr bit of one indicates that those commands still queued after the device has entered the contingent allegiance condition shall be aborted when that condition has terminated. A unit attention condition is created for each initiator that had commands in the queue, but not for the initiator detecting the original error.
- [5] A Disable Queuing (DQue) bit of one indicates that tagged queuing is disabled on the drive. Any pending commands in the queue for that I T X nexus is aborted. Any subsequent queue tag message received shall be rejected with a Message Reject message and the I/O process shall be executed as an untagged command. A DQue bit of zero indicates that tagged queuing is enabled, if the drive supports tagged Queuing.
- [6] An Enable Extended Contingent Allegiance (EECA) bit of one indicates that the device has enabled the extension of the contingent allegiance condition. An EECA bit of zero indicates that the extension of the contingent allegiance condition is disabled.
- [7] See drive Product Manual Mode Sense Data section for changeable values.
- [8] The Busy Timeout Period field specifies the maximum time, in 100 millisecond increments, that the initiator allows for the target to remain busy for unanticipated conditions that are not a routine part of commands from the initiator. This value may be rounded down as defined in section 4.8 of this manual. A 0000h value in this field is undefined by this specification. An FFFFh in this field is defined as an unlimited period.
- [9] GLTSD: Global Logging target save disable bit of zero allows the target to provide a target-defined method for saving log parameters (Logged to disk). A GLTSD bit of one indicates that either the target has disabled the target-defined method for saving log parameters or when set by the initiator specifies that the target-defined method shall be disabled.
- [10] RAC: Report a check. Not used at this time.

Notch Page (0Ch)

The notch page (Table 5.2.1-30) contains parameters for direct access devices that implement a variable number of blocks per cylinder and support this page. Each section of the drive with a different number of blocks per cylinder is referred to as a notch.

Table 5.2.1-30. Notch Page

Bit Byte(s)	7	6	5	4	3	2	1	0
0	PS [1]	Rsvd 0	Page Code (0Ch) 0 0 1 1 0 0					
1	Page Length (16h)							
2 DEFAULT CHANGEABLE	ND [2]	LPN [3]	Reserved 0 0 0 0 0 0					
3 DEFAULT CHANGEABLE	0 0		Reserved 0 0 0 0 0 0					
4,5 DEFAULT CHANGEABLE	(MSB)		Maximum Number of Notches				[4] (LSB)	
6,7 DEFAULT CHANGEABLE	(MSB)		Active Notch				[5] (LSB)	
8-11 DEFAULT CHANGEABLE	(MSB)		Starting Boundary				[6] (LSB)	
12-15 DEFAULT CHANGEABLE	(MSB)		Ending Boundary				[7] (LSB)	
16-23 DEFAULT CHANGEABLE	(MSB)		Pages Notched				[8] (LSB)	

Notes.

- [1] The parameters savable (PS) bit is only used with the Mode Sense command. This bit is reserved with the Mode Select command. A PS bit of one indicates that the drive is capable of saving the page in a non-volatile vendor-specific location.
- [2] A notched drive (ND) bit of zero indicates that the device is not notched and that all other parameters in this page shall be returned as zero by the drive. A ND bit of one indicates that the drive is notched. For each supported active notch value this page defines the starting and ending boundaries of the notch.
- [3] A logical or physical notch (LPN) bit of zero indicates that the notch boundaries are based on the physical parameters of the drive. The cylinder is considered most significant, the head least significant. A LPN bit of one indicates that the notch boundaries are based on logical blocks on the drive.
- [4] The maximum number of notches field indicates the maximum number of notches supported by the drive.

continued from previous page

- [5] The active notch field indicates the notch to which this and subsequent Mode Select and Mode Sense commands shall refer, until the active notch is changed by a later Mode Select command. The value of the active notch shall be greater than or equal to 0 and less than or equal to the maximum number of notches. A active notch value of zero indicates that this and subsequent Mode Select and Mode Sense commands refer to the parameters that apply across all notches.
- [6] The starting boundary field indicates the beginning of the active notch or, if the active notch is zero, the beginning boundary of the drive. If the LPN bit is one, then the four bytes represent a logical block address. If the LPN bit is zero, then the three most significant bytes shall represent the cylinder number and the least significant byte shall represent the head number. When used with the Mode Select command this field is ignored.
- [7] The ending boundary field indicates the ending of the active notch or, if the active notch is zero, the ending of the drive. If the LPN bit is one, then the four bytes represent logical block address. If the LPN bit is zero, then the three most significant bytes shall represent the cylinder number and the least significant byte shall represent the head number. When used with the Mode Select command this field is ignored.

Each notch shall span a set of consecutive logical blocks on the drive, the notches shall not overlap, and no logical block shall be excluded from a notch.

- [8] The pages notched field is a bit map of the mode page codes that indicates which pages contain parameters that may be different for different notches. The most significant bit of this field corresponds to page code 3Fh and the least significant bit corresponds to page code 00h. If a bit is one, then the corresponding mode page contains parameters that may be different for different notches. If a bit is zero, then the corresponding mode page contains parameters that are constant for all notches.
- [9] See individual drive Product Manual Mode Sense Data section for a table showing codes that indicate which of these bits are changeable by the host using the Mode Select command.
- [10] See drive Product Manual Mode Sense Data section for changeable values.

Table 5.2.1-31. Power Condition Page

Bit Byte(s)	7	6	5	4	3	2	1	0
0	PS	Reserved		Page Code [5]				
1	Page Length (0Ah)							
2	Reserved							
3	Reserved						Idle [1]	Standby [2]
4	(MSB) Idle Condition Timer [3] (LSB)							
5								
6								
7								
CHANGEABLE	[6]							
8	(MSB) Standby Condition Timer [4] (LSB)							
9								
10								
11								
CHANGEABLE	[6]							

The power condition page provides the initiator the means to control the length of time a logical unit will delay before changing its power requirements. There is no notification to the initiator that a logical unit has entered into one of the power conditions.

On the receipt of a command the device shall adjust itself to the power condition which allows the command to execute. The timer which maps to this power condition and any lower power condition timers shall be reset on receipt of the command. The timer associated with this power condition shall be restarted when the condition that forces the change in power completes.

Notes [].

- [1] An idle bit of one indicates a logical unit shall use the Idle condition Timer to determine the length of inactivity time to wait before entering the Idle condition. An idle bit of zero indicates a logical unit shall not enter the Idle condition.
- [2] A standby bit will not be supported at this time.
- [3] The Idle Condition Timer field indicates the inactivity time in 100 millisecond increments that the logical unit shall wait before entering the Idle condition. Minimum time is 500 milliseconds. The Idle Condition Timer field is not supported at this time.
- [4] The Standby condition Timer field is not supported at this time.
- [5] Some drive models use page code 0Dh, others use 1Ah. See drive Product Manual Vol. 1, Mode Sense Data and Commands supported Tables for indication of which the drive model in question uses.
- [6] See individual drive Product Manual Mode Sense Data section for a table showing codes that indicate which of these bits are changeable by the host using the Mode Select command.

Xor Control Mode page (10h)

The xor control mode page (see Table 5.2.1-32) provides the initiator with the means to obtain or modify certain xor operating parameters of the target.

Table 5.2.1-32. Xor Control Mode page (10h)

Bit Byte(s)	7	6	5	4	3	2	1	0
0	PS	Reserved	Page Code (10h)					
1	Page Length (16h)							
2	Reserved						XorDis[1]	Reserved
3	Reserved							
4	MSB Maximum xor write size[2] LSB							
5								
6								
7								
CHANGEABLE	[6]							
8	MSB Reserved LSB							
9								
10								
11								
CHANGEABLE	[6]							
12	MSB Maximum regenerate size [3] LSB							
13								
14								
15								
CHANGEABLE	[6]							
16	MSB Maximum rebuild read size [4] LSB							
17								
18								
19								
CHANGEABLE	[6]							
20	Reserved							
21	Reserved							
CHANGEABLE	[6]							
22	Rebuild delay [5]							
23	Rebuild delay [5]							
CHANGEABLE	[6]							

Notes.

- [1] An xor disable (XorDis) bit of zero enables the xor operations within a device. An XorDis bit of one disables the xor operations within a device. If the XorDis bit is set to one and an xor command is sent to the target the command shall be terminated with CHECK CONDITION status. The sense data shall be set to ILLEGAL REQUEST: INVALID COMMAND OPERATION CODE.
- [2] The maximum xor write size field specifies the maximum transfer length in blocks that the target accepts for a single XDWRITE EXTENDED, XDWRITE, or XPWRITE command.
- [3] The maximum regenerate size field specifies the maximum regenerate length in blocks that the target accepts for the REGENERATE command.
- [4] The maximum rebuild read size field specifies the maximum rebuild length in blocks that the target shall use for READ commands during a rebuild operation. This field does not limit the rebuild size.
- [5] The rebuild delay field specifies the minimum time in milliseconds between successive READ commands during a rebuild operation.
- [6] See individual drive Product Manual Mode Sense Data section for a table showing codes that indicate which of these bits are changeable by the host using the Mode Select command.

Table 5.2.1-33. Unit Attention Parameters (00h)

Bit Byte(s)	7	6	5	4	3	2	1	0
0	PS 1[10]	0	0	0	0	0	0	0
1	Page Length (in bytes) [1]							
2	0	SSM [2]	IL [3]	Unit- Attn[4]	0	Rnd [5]	Strict [6]	S2 [7]
CHANGEABLE	[9]							
3	[8]							
CHANGEABLE	[9]							

Page zero is the last page to be reported by the drive.

Notes.

- [1] The page length field specifies the length in bytes of the mode parameters that follow. If the initiator does not set this value to the value that is returned for the page by the MODE SENSE command, the drive shall terminate the command with CHECK CONDITION status. The sense key shall be set to ILLEGAL REQUEST with the additional sense code set to INVALID FIELD IN PARAMETER LIST. The drive is permitted to implement a mode page that is less than the full page length defined by this specification, provided no field is truncated and the page length field correctly specifies the actual length implemented. *If the Strict bit equals zero and if the page length specified by the initiator is shorter than the actual page length, then the parameters are transferred and the command ends with GOOD status if no other items cause the command to be rejected. **Caution: Utilization of this forgiving option by an initiator that does not analyze the impact of the truncation could adversely affect data integrity.***
- [2] If the enable Synchronous Select Mode (SSM) bit equals one, the drive initiates WDTR and SDTR messages when it recognizes that one may be required (after reset, reset msg. or power cycle). If the SSM bit equals zero the drive does not initiate WDTR or SDTR regardless of negotiated conditions prior to reset, reset message or power cycle.
- [3] "Inquiry Length" (IL) bit. When the IL bit is set to 1, the standard INQUIRY data available to a host is limited to the 36 bytes required by the SCSI-2 specification. When the IL bit is reset (0), 148 bytes of standard INQUIRY data are available. The Additional Length field in byte 4 of the INQUIRY data is updated to reflect the actual number of additional bytes available.
- [4] When the Unit Attn bit is set to 1, then Unit Attention is logged in sense only; no Check Condition Status is presented following any reset. When this bit is RESET (0) then Check Condition is presented for all affected Initiators following a reset until Request Sense is issued by each Initiator (as per current operation).
- [5] When the Round bit equals one, the drive treats and reports rounded parameters as described in Section 4.8. When Round equals zero the drive shall round the parameter and handle command completion reporting as if the parameter had not been rounded.
- [6] When the Strict bit is a one, the drive checks for Initiator attempts to change unchangeable parameters. If the drive detects an attempt it rejects the command in the standard way, i.e., Check Condition status from drive, Request Sense from the Initiator, Illegal Request sense key (5h) back from the drive. When the Strict bit is zero, the drive ignores the values of the unchangeable parameters in a Mode Select command. The drive does not reject the command trying to change unchangeable parameters.

continued from previous page

[7] "SCSI-2" (S2) bit

When set to one the SCSI-2 (S2) bit changes the following SCSI-3 features from their SCSI-3 definition to the SCSI-2 definition. When S2 equals zero, the following features remain as specified in other portions of this specification.

- a) Control Mode Page (0Ah) Length from 0Ah to 06h.
- b) Caching Page (08) Length from 12h to 0Ah.

[8] This byte is reserved for future compatibility with Seagate host adapters. Though presently may be changeable, (see note [7]) this byte does not control anything, unless the individual drive Product Manual indicates that it does and defines its use in the Mode Sense Data section.

[9] See individual drive Product Manual Mode Sense Data section for a table showing codes that indicate which of these bits are changeable by the host using the Mode Select command.

[10] A PS bit of one indicates that the drive is capable of saving the page in a nonvolatile vendor-specific location (used only with Mode Sense command).

Informational Exceptions Control page (1Ch)

The informational exceptions control page (see Table 5.2.1-33a) defines the methods used by the target to control the reporting and the operations of specific informational exception conditions. This page shall only apply 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 target. An informational exception condition may occur asynchronous to any commands issued by an application client.

Mode page 1Ch may be used by the drive to implement the S.M.A.R.T. system. S.M.A.R.T. is an acronym for Self-Monitoring Analysis and Reporting Technology. The intent of the S.M.A.R.T. system is to recognize conditions that indicate imminent drive failure, and provide sufficient warning to the host system of impending failure. Thus the host system may perform data backup before it is too late.

Table 5.2.1-33a. Informational Exceptions Control Page (1Ch)

PAGE DESCRIPTOR HEADER								
Bit Byte(s)	7	6	5	4	3	2	1	0
0	PS	Reserved	Page code (1Ch)					
1	Page Length (0Ah)							

INFORMATIONAL EXCEPTIONS CONTROL PARAMETERS

2	Perf [1]	Reserved	DExcpt[2]	Test[3]	Reserved	LogErr[4]
CHANGEABLE	[9]					
3	Reserved	MRIE [5]				
CHANGEABLE	[9]					
4	(MSB)					
5	Interval timer [6] [8]					
6						
7						(LSB)
CHANGEABLE	[9]					
8	(MSB)					
9	Report count [7] [8]					
10						
11						(LSB)
CHANGEABLE	[9]					

Notes.

- [1] A Performance bit (Perf) of zero indicates that informational exception operations that are the cause of delays are acceptable. A Perf bit of one indicates the target shall not cause delays while doing informational exception operations. A Perf bit set to one may cause the target to disable some or all of the informational exception operations, thereby limiting the reporting of informational exception conditions.
- [2] A disable exception control (DExcpt) bit of zero indicates information exception operations shall be enabled. The reporting of information exception conditions when the DExcpt bit is set to zero is determined from the method of reporting informational exceptions field. A DExcpt bit of one indicates the target shall disable all information exception operations. The method of reporting informational exceptions field is ignored when DExcpt is set to one.
- [3] A test bit (Test) of one instructs the drive to create false drive failure(s) at the next interval time, provided that the DExcpt bit is not set. The MRIE and Report Count fields apply as specified in this document. A false drive failure will be reported as sense code/qualifier 5DFF (FF for false failure versus a true failure 5D00). A Test bit of zero instructs the drive to not generate any false drive failure notifications.

-
- [4] The log errors bit (LogErr) of zero indicates that the logging of informational exception conditions within a target is vendor specific. A LogErr bit of one indicates the target shall log informational exception conditions.
- [5] The Method of Reporting Informational Exceptions field (MRIE) indicates the methods that shall be used by the target to report informational exception conditions (see Table 5.2.1-33b). The priority of reporting multiple information exceptions is vendor specific.
- [6] The Interval Timer field indicates the period in 100 millisecond increments for reporting that an informational exception condition has occurred. The target shall not report informational exception conditions more frequently than the time specified by the Interval Timer field and as soon as possible after the timer interval has elapsed. After the informational exception condition has been reported the interval timer shall be restarted. A value of zero in the Interval Timer field indicates that the target shall only report the informational exception condition one time. A value of FFFFFFFFh in the Interval Timer field shall indicate the timer interval is vendor specific. If the method of reporting informational exceptions field has a value of six then the Interval Timer field is ignored.
- [7] The Report Count field indicates the number of times to report an informational exception condition to the application client. A value of zero in the Report Count field indicates there is no limit on the number of times the target shall report an informational exception condition.
- [8] The maintaining of the Interval Timer and the Report Count after power cycles and/or resets by the target shall be vendor specific. [Seagate: A predicted failure is preserved through resets and power cycles. The count of how many informational exception conditions that were reported by the target to any initiator is cleared by a reset of power cycle.]
- [9] See individual drive Product Manual Mode Sense Data section for a table showing codes that indicate which of these bits are changeable by the host using the Mode Select command.

Table 5.2.1-33b. Format of REPORTING METHOD Field

Code	Description
0h	<p>No reporting of informational exception condition: This method instructs the target to not report information exception conditions.</p>
1h	<p>Asynchronous event reporting: This method instructs the target to report informational exception conditions by using the rules for asynchronous event reporting as described in the SCSI-3 Architectural Model and the relevant Protocol Standard.</p> <p>The sense key shall be set to RECOVERED ERROR and the additional sense code shall indicate the cause of the informational exception condition.</p>
2h	<p>Generate Unit Attention: This method instructs the target to report informational exception conditions by returning a CHECK CONDITION status on any command. The sense key shall be set to UNIT ATTENTION and the additional sense code shall indicate the cause of the informational exception condition.</p> <p>The command that has the CHECK CONDITION shall not be executed before the informational exception condition is reported.</p>
3h	<p>Conditionally generate recovered error: This method instructs the target to report informational exception conditions, dependent on the value of the PER bit of the error recovery parameters mode page, by returning a CHECK CONDITION status on any command. The sense key shall be set to RECOVERED ERROR and the additional sense code shall indicate the cause of the informational exception condition.</p> <p>The command that has the CHECK CONDITION shall complete without error before any informational exception condition may be reported.</p>
4h	<p>Unconditionally generate recovered error: This method instructs the target to report informational exception conditions, regardless of the value of the PER bit of the error recovery parameters mode page, by returning a CHECK CONDITION status on any command. The sense key shall be set to RECOVERED ERROR and the additional sense code shall indicate the cause of the informational exception condition.</p> <p>The command that has the CHECK CONDITION shall complete without error before any informational exception condition may be reported.</p>
5h	<p>Generate no sense: This method instructs the target to report informational exception conditions by returning a CHECK CONDITION status on any command. The sense key shall be set to NO SENSE and the additional sense code shall indicate the cause of the informational exception condition.</p> <p>The command that has the CHECK CONDITION shall complete without error before any informational exception condition may be reported.</p>
6h	<p>Only report information exception condition on request: This method instructs the target to preserve the informational exception(s) information. To find out about information exception conditions the Application Client polls the target by issuing an unsolicited REQUEST SENSE command. The sense key shall be set to NO SENSE and the additional sense code shall indicate the cause of the informational exception condition.</p>
7h - Bh	Reserved
Ch - Fh	Vendor Specific

5.2.1.11 Start Stop Unit Command (1Bh)

Table 5.2.1-34. Start/Stop Unit Command

Bit Byte(s)	7	6	5	4	3	2	1	0
0	0	0	0	1	1	0	1	1
1	Logical Unit No. 0 0 0 0 0 0						0	IMMED [1]
2	0	0	0	0	0	0	0	0
3	0	0	0	0	0	0	0	0
4	0	0	0	0	0	0	0	Start [2]
5	0	0	0	0	0	0	Flag	Link [3]

The Start/Stop Unit command requests that the target enable the logical unit for further operations (start), or stop spindle rotation (stop).

Notes.

- [1] An Immed bit of zero indicates that status shall be returned after the operation is completed. If the Immed bit is a one status is returned as soon as the operation is initiated.
- [2] When the Start bit is a one, this requests that the logical unit be made ready for use. If the Start bit is zero (requesting that the unit be stopped), the request causes the drive to either actually stop the spindle or to simulate the drive spindle stopped condition. In this latter situation the drive spindle actually continues to spin and the drive reports Not Ready in response to media access commands. Power consumption is as when in idle mode. The type of stop implemented is given in the drive Product Manual.
- [3] See "Control Byte" paragraph 4.2.6.

The drive is able to execute the following commands when the drive spindle is not rotating, or in a simulated stopped condition.

Test Unit Ready	Motor Start
Request Sense	Receive Diagnostic Results
Inquiry	Write Data Buffer
Reserve	Read Data Buffer
Release	

The remaining commands (see section "SCSI Interface commands supported" in drive Product Manual for list of all commands supported) cannot be executed until after the drive has spindled up. If the drive receives one of these commands before it can be executed, a "Check Condition" status is returned (with Sense Key of "Not Ready").

For systems that support disconnection, the drive disconnects when a Start Unit procedure is commanded, and reconnects when the unit is up to speed and Ready or when operation is initiated if the "Immed" bit is one.

5.2.1.12 Prevent/Allow Medium Removal Command (1Eh)

Not implemented by the drive. If received, the drive terminates with Check Condition status and set an Illegal Request Sense Key.

5.2.2 Commands for Direct Access Devices

The commands that may be implemented by the drive are listed in Table 5.2.2-1.

See section “SCSI Interface commands supported” in the individual drive Product Manuals for complete list of commands supported.

Table 5.2.2-1. Commands for Direct Access Devices

Operation Code (Hex)	Command Name	Section	Page
25	READ CAPACITY	5.2.2.1	184
28	READ EXTENDED	5.2.2.2	186
2A	WRITE EXTENDED	5.2.2.3	188
2B	SEEK EXTENDED	5.2.2.4	189
2E	WRITE AND VERIFY	5.2.2.5	190
2F	VERIFY	5.2.2.6	191
34	PREFETCH	5.2.2.16	203
35	SYNCHRONIZE CACHE	5.2.2.7	192
37	READ DEFECT DATA	5.2.2.8	193
3B	WRITE DATA BUFFER	5.1.2.3	108
3C	READ DATA BUFFER	5.1.2.4	113
3E	READ LONG	5.2.2.9	195
3F	WRITE LONG	5.2.2.10	197
41	WRITE SAME	5.2.2.11	198
56	RESERVE (10)	5.2.2.12	199
57	RELEASE (10)	5.2.2.13	200
55	MODE SELECT (10)	5.2.2.14	201
5A	MODE SENSE (10)	5.2.2.15	202

Operation Codes 50, 51, 52, 80, 81, 82 are described in Section 5.3.

5.2.2.1 Read Capacity Command (25H)

Table 5.2.2-2. Read Capacity Command

Bit Byte(s)	7	6	5	4	3	2	1	0
0	0	0	1	0	0	1	0	1
1	Logical Unit No. [1]						RelAdr [2]	
2	Logical Block Address (MSB)						[3]	
3	Logical Block Address						[3]	
4	Logical Block Address						[3]	
5	Logical Block Address (LSB)						[3]	
6	0	0	0	0	0	0	0	0
7	0	0	0	0	0	0	0	0
8	0	0	0	0	0	0	0	PMI [4]
9	0	0	0	0	0	0	Flag	Link [5]

The Read Capacity command (Table 5.2.2-2) provides a means for the initiator to request the capacity of the drive information.

Notes.

- [1] LUN must be zero.
- [2][3] A relative address (RelAdr) bit of one indicates that the logical block address field [3] is a two's complement displacement. This negative or positive displacement is to be added to the logical block address last accessed on the logical unit to form the logical block address [3] for this command. This feature is only available when linking commands. The feature requires that a previous command in the linked group have accessed a block of data on the logical unit.
- A RelAdr bit of zero indicates that the logical block address field [3] specifies the first logical block of the range of logical blocks to be operated on by this command.
- [4] A Partial Medium Indicator (PMI) bit of zero indicates the information returned in the Read Capacity data shall be the Logical Block Address and Block Length (in bytes) of the last logical block [6] of the logical unit. The Logical Block Address in the Command Descriptor Block shall be set to zero for this option.
- A PMI bit of one indicates the information returned shall be the Logical Block Address and Block Length (in bytes) of the last Logical Block Address after which a substantial delay (defined as approximately 1 millisecond for the typical drive) in data transfer is encountered. This returned Logical Block Address shall be greater than or equal to the Logical Block Address specified in the Command Descriptor Block. This reported Logical Block Address is the last block prior to a cylinder boundary.
- [5] See "Control Byte" paragraph 4.2.6.
- [6] This value is the same for all drives of the same model number, sector size and sparing scheme. It is the same regardless of the number of defective blocks the drive has.

In response to the Read Capacity Command, the drive returns 8 bytes of Read Capacity Data to the host. The contents of the 8 bytes are listed in Table 5.2.2-2a.

Table 5.2.2-2a. Read Capacity Data

Byte	Description
0	LOGICAL BLOCK ADDRESS (MSB)
1	LOGICAL BLOCK ADDRESS
2	LOGICAL BLOCK ADDRESS
3	LOGICAL BLOCK ADDRESS (LSB)
4	BLOCK LENGTH (MSB)
5	BLOCK LENGTH
6	BLOCK LENGTH
7	BLOCK LENGTH (LSB)

5.2.2.2 Read Extended Command (28h)

Table 5.2.2-3. Read Extended Command

Bit Byte(s)	7	6	5	4	3	2	1	0
0	0	0	1	0	1	0	0	0
1	Logical Unit No.[1] 0 0		0	DPO [2]	FUA [3]	0	0	RelAdr [4]
2	Logical Block Address (MSB)							[5]
3	Logical Block Address							[5]
4	Logical Block Address							[5]
5	Logical Block Address (LSB)							[5]
6	0	0	0	0	0	0	0	0
7	Transfer Length (MSB)							[6]
8	Transfer Length (LSB)							[6]
9	0	0	0	0	0	0	Flag [7]	Link [7]

The Read Extended Command requests that the target transfer data to the initiator. This command is implemented with the drive specific parameters shown in Table 5.2.2-3.

This command operates the same as the Read command (see Section 5.2.1.4) except that in the CDB for this command a four byte Logical Block Address and a two byte Transfer Length may be specified.

The data value most recently written in the addressed logical block is returned to the Host.

This command shall be terminated with a Reservation Conflict status if any reservation access conflict (see Section 5.2.1.8) exists, and no data shall be read.

Notes.

[1] LUN must be zero.

[2] A disable page out (DPO) bit of one indicates that the drive shall assign the logical blocks accessed by this command the lowest priority for being fetched into or retained by the cache.

The DPO bit is used to control replacement of logical blocks in the cache memory when the host has information on the future usage of the logical blocks. If the DPO bit is set to one, the host knows the logical blocks accessed by the command are not likely to be accessed again in the near future and should not be put in the cache memory nor retained by the cache memory. If the DPO bit is zero, the host expects that logical blocks accessed by this command are likely to be accessed again in the near future.

[3] A force unit access (FUA) bit of one indicates that the target shall access the media in performing the command prior to returning GOOD status. Read commands shall access the specified logical blocks from the media (i.e., the data is not directly retrieved from the cache). In the case where the cache contains a more recent version of a logical block than the media, the logical block shall first be written to the media.

continued from previous page

An FUA bit of zero indicates that the target may satisfy the command by accessing the cache memory. For read operations, any logical blocks that are contained in the cache memory may be transferred to the initiator directly from the cache memory.

- [4] A relative address (RelAdr) bit of one indicates that the logical block address field is a two's complement displacement. This negative or positive displacement is to be added to the logical block address last accessed on the logical unit to form the logical block address for this command. This feature is only available when linking commands. The feature requires that a previous command in the linked group have accessed a block of data on the logical unit.

A RelAdr bit of zero indicates that the logical block address field specifies the first logical block of the range of logical blocks to be operated on by this command.

- [5] The Logical Block Address specifies the logical block at which the read operation shall begin, if RelAdr bit is zero (see description of RelAdr bit following).

- [6] The Transfer Length specifies the number of contiguous logical blocks of data that shall be transferred. A Transfer Length of zero indicates that no logical blocks shall be transferred. This condition shall not be considered an error. Any other value indicates the number of logical blocks that shall be transferred.

- [7] See "Control Byte" paragraph 4.2.6.

If any of the following conditions occur, this command shall return a Check Condition status and the Sense Key shall be set as indicated. This table does not provide an exhaustive enumeration of all conditions that may cause the Check Condition status.

Condition	Sense Key
Invalid Logical Block Address	Illegal Request (see note)
Target reset since last command from this initiator	Unit Attention
Unrecovered read error	Medium Error
Recoverable read error	Recovered Error

Note. The extended sense information bytes shall be set to the Logical Block Address of the first invalid address.

5.2.2.3 Write Extended Command (2Ah)

Table 5.2.2-4. Write Extended Command

Bit Byte(s)	7	6	5	4	3	2	1	0
0	0	0	1	0	1	0	1	0
1	Logical Unit No.[1] 0 0		0	DPO [2]	FUA [3]	0	0	RelAdr [4]
2	Logical Block Address (MSB)							[5]
3	Logical Block Address							[5]
4	Logical Block Address							[5]
5	Logical Block Address (LSB)							[5]
6	0	0	0	0	0	0	0	0
7	Transfer Length (MSB)							[6]
8	Transfer Length (LSB)							[6]
9	0	0	0	0	0	0	Flag	Link [7]

The Write Extended command requests that the drive write to the medium the data transferred by the initiator.

This command is implemented with the drive specific parameters listed in Table 5.2.2-4. Refer also to "Write Caching" section in the individual Product Manual for information on Write Cache Control.

Notes.

- [1] In the CDB, the LUN must be zero.
- [2] If the Disable Page Out (DPO) bit is set to one, no data is cached. The DPO bit is only meaningful if the RCD bit of Mode Select Page 8 is set false (Caching enabled).
- [3] A force unit access (FUA) bit of one indicates that the write command shall not return GOOD status until the logical blocks have actually been written on the media. The FUA bit is only meaningful if the WCE bit of Mode Sense page 8 is true.
- [4] A relative address (RelAdr) bit of one indicates that the logical block address field is a two's complement displacement. This negative or positive displacement is to be added to the logical block address last accessed on the logical unit to form the logical block address for this command. This feature is only available when linking commands. The feature requires that a previous command in the linked group have accessed a block of data on the logical unit.

A RelAdr bit of zero indicates that the logical block address field specifies the first logical block of the range of logical blocks to be operated on by this command.

- [5] The Logical Block Address specifies the logical block at which the write operation shall begin, if RelAdr bit is zero. (See description of RelAdr bit).
- [6] The Transfer Length specifies the number of contiguous logical blocks of data that shall be transferred. A Transfer Length of zero indicates no logical blocks shall be transferred. This condition shall not be considered an error and no data shall be written. Any other value indicates the number of logical blocks that shall be transferred.

- [7] See "Control Byte" paragraph 4.2.6.

This command operates the same as the Write command (Section 5.2.1.5) except that in the CDB for this command a four byte Logical Block Address and a two byte Transfer Length may be specified.

This command shall be terminated with a Reservation Conflict status if any reservation access conflict (see Section 5.2.1.8) exists, and no data shall be written.

If any of the following conditions occur, this command shall be terminated with a Check Condition status and the Sense Key shall be set as indicated in the following table. This table does not provide an exhaustive enumeration of all conditions that may cause the Check Condition status.

Condition	Sense Key
Invalid Logical Block Address	Volume Overflow (see note)
Target reset since the last command from this initiator	Unit Attention

Note. The extended sense information bytes shall be set to the Logical Block Address of the first invalid address.

5.2.2.4 Seek Extended Command (2Bh)

Table 5.2.2-5a. Seek Extended Command

Bit Byte(s)	7	6	5	4	3	2	1	0
0	0	0	1	0	1	0	1	1
1	Logical Unit No.[1] 0 0 0			0	0	0	0	0
2	Logical Block Address (MSB)							[2]
3	Logical Block Address							[2]
4	Logical Block Address							[2]
5	Logical Block Address (LSB)							[2]
6	0	0	0	0	0	0	0	0
7	0	0	0	0	0	0	0	0
8	0	0	0	0	0	0	0	0
9	0	0	0	0	0	0	Flag	Link [3]

The Seek Extended command requests that the drive seek to the specified Logical Block Address.

This command is implemented with the drive specific parameters listed in Table 5.2.2-5a

Notes.

[1] In the CDB, the LUN must be zero.

[2] This command operates the same as the Seek command (Section 5.2.1.6) except that a four byte Logical Block Address is specified.

[3] See "Control Byte" paragraph 4.2.6.

5.2.2.5 Write and Verify Command (2Eh)

Table 5.2.2-5b. Write and Verify Command

Bit Byte(s)	7	6	5	4	3	2	1	0
0	0	0	1	0	1	1	1	0
1	Logical Unit No.[1] 0 0		0	DPO 0 [2]	0	0	BYT CHK [3]	RelAdr [4]
2	Logical Block Address (MSB)							[5]
3	Logical Block Address							
4	Logical Block Address							
5	Logical Block Address (LSB)							
6	0	0	0	0	0	0	0	0
7	Transfer Length (MSB)							[6]
8	Transfer Length (LSB)							
9	0	0	0	0	0	0	Flag	Link [7]

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

Notes.

- [1] The LUN must be zero.
 [2] A disable page out (DPO) bit of one indicates that the target shall assign the logical blocks accessed by this command the lowest priority for being fetched into or retained by the cache.

The DPO bit is used to control replacement of logical blocks in the cache memory when the host has information on the future usage of the logical blocks. If the DPO bit is set to one, the logical blocks accessed by the command are not likely to be accessed again in the near future and should not be put in the cache memory nor retained by the cache memory. If the DPO bit is zero, that logical blocks accessed by this command are likely to be accessed again in the near future.

- [3] A byte check (BytChk) bit of zero causes the verification to be simply a medium verification (ECC) with no data comparison. A BytChk bit of one causes a byte-by-byte compare of data written on the peripheral device and the data transferred from the initiator. If the compare is unsuccessful, the command shall be terminated with a CHECK CONDITION status and the sense key shall be set to MISCOMPARE.
 [4] RelAdr. See Description of RelAdr bit for the Write Extended Command.
 [5] The logical block address specifies the logical block at which the write operation shall begin if RelAdr bit is zero. (see description of RelAdr bit for the Write Extended Command).
 [6] The transfer length specifies the number of contiguous logical blocks of data that shall be transferred. A transfer length of zero indicates that no logical blocks shall be transferred. This condition shall not be considered as an error and no data shall be written. Any other value indicates the number of logical blocks that shall be transferred.

For Systems that support disconnection, the drive disconnects during the execution of this command.

- [7] See "Control Byte" paragraph 4.2.6.

5.2.2.6 Verify Command (2Fh)

Table 5.2.2-6. Verify Command

Bit Byte(s)	7	6	5	4	3	2	1	0	
0	0	0	1	0	1	1	1	1	
1	Logical Unit No. [1] 0 0 0			DPO 0 [2]		0	0	BYT CHK[3]	RelAdr [4]
2	Logical Block Address (MSB)							[5]	
3	Logical Block Address							[5]	
4	Logical Block Address							[5]	
5	Logical Block Address (LSB)							[5]	
6	0	0	0	0	0	0	0	0	
7	Verification Length (MSB)							[6]	
8	Verification Length (LSB)							[6]	
9	0	0	0	0	0	0	Flag	Link [7]	

The Verify command requests that the target verify the data written on the medium.

This command is implemented with the drive specific parameters listed in Table 5.2.2-6.

The drive disconnects while this command is being executed if the initiator supports disconnect/reconnect.

Notes.

- [1] LUN must be zero.
 - [2] A disable page out (DPO) bit of one indicates that the target shall assign the logical blocks accessed by this command the lowest priority for being fetched into or retained by the cache.
- The DPO bit is used to control replacement of logical blocks in the cache memory when the host has information on the future usage of the logical blocks. If the DPO bit is set to one, the logical blocks accessed by the command are not likely to be accessed again in the near future and should not be put in the cache memory nor retained by the cache memory. If the DPO bit is zero, that logical blocks accessed by this command are likely to be accessed again in the near future.
- [3] A Byte Check (BytChk) bit of zero causes the verification to be simply a medium verification (CRC, ECC, etc). A BytChk bit of one causes a byte by byte compare of data on the medium and the data transferred from the initiator. If the compare is unsuccessful, the command shall be terminated with a Check Condition status and the Sense Key shall be set to Miscompare.
 - [4] RelAdr: See description of RelAdr bit for Write Extended command.
 - [5] The Logical Block Address specifies the logical block at which the verify operation shall begin, if RelAdr is Zero.
 - [6] The Verification Length specifies the number of contiguous logical blocks of data that shall be verified. A Verification Length of zero indicates that no logical blocks shall be verified (an Implied Seek is still performed). This condition shall not be considered an error. Any other value indicates the number of logical blocks that shall be verified.
 - [7] See "Control Byte" paragraph 4.2.6.

5.2.2.7 Synchronize Cache Command (35h)

Table 5.2.2-7. Synchronize Cache Command

Bit Byte(s)	7	6	5	4	3	2	1	0
0	0	0	1	1	0	1	0	1
1	Logical Unit No. [1]			Reserved			Immed [2]	RelAdr [3]
2 : 5	(MSB) Logical Block Address							[4] (LSB)
6	Reserved							
7 : 8	(MSB) Number of Blocks							[5] (LSB)
9	Control							[6]

The Synchronize Cache command (Table 5.2.2-7) ensures that logical blocks in the cache memory, within the specified range, have their most recent data value recorded on the physical medium. If a more recent data value for a logical block within the specified range exists in the cache memory than on the physical medium, then the logical block from the cache memory shall be written to the physical medium. Logical blocks are not necessarily removed from the cache memory as a result of the synchronize cache operation.

[1] LUN must be zero.

[2] An immediate (Immed) bit of one indicates that the drive shall return status as soon as the command descriptor block has been validated. An Immed bit of zero indicates that the status shall not be returned until the operation has been completed. If the Immed bit is one and the target does not support it, the command shall terminate with CHECK CONDITION status. The sense key shall be set to ILLEGAL REQUEST and the additional sense code shall be set to INVALID FIELD IN CDB.

[3] See paragraph 5.2.2.9 note [3] for definition of RELADR bit.

[4] The Logical block address field specifies the logical block at which the Synchronize Cache operation begins.

[5] The number of blocks field specifies the total number of contiguous logical blocks within the range. A number of blocks of zero indicates that all remaining logical blocks on the logical unit shall be within the range.

A logical block within the specified range that is not in cache memory is not considered an error. Multiple locks may be in effect from more than one initiator. Locks from different initiators may overlap. An unlock of an overlapped area does not release the lock of another initiator.

[6] See section 4.2.6.

5.2.2.8 Read Defect Data Command (37h)

Table 5.2.2-8a. Read Defect Data Command

Bit Byte(s)	7	6	5	4	3	2	1	0
0	0	0	1	1	0	1	1	1
1	Logical Unit No.[1] 0 0 0			0	0	0	0	0
2	0	0	0	P	G	Defect List Format 1 0 0 [3] or 1 0 1		
3	0	0	0	0	0	0	0	0
4	0	0	0	0	0	0	0	0
5	0	0	0	0	0	0	0	0
6	0	0	0	0	0	0	0	0
7	Allocation Length (MSB)							[4]
8	Allocation Length (LSB)							[4]
9	0	0	0	0	0	0	Flag	Link[5]

The Read Defect Data command requests that the target transfer the medium defect data to the initiator. If the drive is unable to access any medium defect data, it terminates the command with Check Condition status. The sense key is set to either Medium Error if a medium error occurred or No Sense if the list does not exist and the additional Sense Code is set to DEFECT LIST NOT FOUND.

This command is implemented with the drive specific parameters listed in Table 5.2.2-8a.

Notes.

[1] LUN must be zero.

[2] The drive interprets the P and G bits (bits 4 and 3 of byte 2 of the CDB) as follows:

Bit P	Bit G	
0	0 (1)	Return Defect List header only
0	1 (2)	Return the growth "G" list only.
1	0 (3)	Return the manufacturers original ETF list only.
1	1 (4)	Return all lists.

(1) Target returns only the defect List header.

(2) This list reflects the grown or "G" list as defined in Section 5.2.1.2.

(3) This list reflects the manufacturers original ETF list. These defects may or may not have been reallocated, depending on the last Format command received (the last format may or may not have requested the P list flaws be reallocated during the format function).

(4) The returned list contains all of the requested drive's defect lists (i.e. P, G, C & D) regardless of whether these lists have been reallocated by the drive.

continued from previous page

- [3] The Defect List format field indicates the defect data format preferred by the initiator. The meaning is the same as the Defect List Format field in the Format command (indicated in Table 5.2.1-3)

The Defect List format bits (bits 2, 1, 0 in the CDB) should be; 1 0 0, respectively, to signify a defect list in the Bytes from Index format, or, 1 0 1, respectively to signify a defect list in the Physical Sector format. If neither of these two, the drive responds with the defect list in the drives default format (physical sector) and create the check condition status with Recovered Error Sense Key [1h] and additional sense error code 1C at the end of the Read Defect Data transfer.

- [4] Allocation Length specifies the number of bytes the initiator has allocated for the returned defect data. An Allocation Length of zero indicates that no Read Defect Data shall be transferred. Any other value indicates the maximum number of bytes to be transferred. The drive shall terminate the Data In phase when the Allocation Length bytes have been transferred or when all available defect data has been transferred to the initiator, whichever is less.

- [5] See "Control Byte" paragraph 4.2.6.

This command is intended to be used only with the Format Unit command (Section 5.2.1.2). The initiator should not interpret or act upon this list except to resend this list as defect data in a Format Unit command. It is not possible to relate actual physical locations to logical block addresses that are given in connection with other commands.

The format Defect Data header and Defect Data Descriptor bytes returned are shown in Table 5.2.2-8b.

The first 4 bytes returned are the Defect List Header. The P bit, G bit, and Defect List Format fields indicate the defect format actually returned by the drive. The definitions are the same as for byte 2 of the Read Defect Data Command Descriptor Block (Table 5.2.2-8a).

The Defect List Length specifies the total length in bytes of all the defect descriptors available from the drive. If the Allocation Length of the CDB is too small to transfer all of the defect descriptors, the Defect List Length is not adjusted to reflect the truncation. The drive does not create the CHECK CONDITION status. The Defect Descriptors are not required to be in ascending order.

Table 5.2.2-8b. Defect List Header Description

Bit Byte(s)	7	6	5	4	3	2	1	0
0	0	0	0	0	0	0	0	0
1	0	Reserved 0	0	P	G	Defect List Format		
				[1]		1 or 1	0 0	0 [2] 1
2	Defect List Length (MSB)							
	(Value is 8 Times the Number of Defects)							
3	Defect List Length (LSB)							
	[3]							
4-n	Defect Descriptor Bytes							
	[3]							

[1] Either 00, 01, 10, or 11 as previously defined in note [1] of Table 5.2.2-8a.

[2] 100 defines a list in the Bytes from Index format.
101 defines a list in the Physical Sector format.

[3] If the P, G bits were 0,0 the Defect List length will be zero and no Defect Descriptor Bytes are sent to the initiator. See Tables 5.2.1-6e and 5.2.1-6f for format of Defect Descriptor Bytes.

5.2.2.9 Read Long Command (3Eh)

Table 5.2.2-9. Read Long Command

Bit Byte(s)	7	6	5	4	3	2	1	0
0	0	0	1	1	1	1	1	0
1	Logical Unit No.[1]						COR- RCT [2]	RelAdr [3]
	0	0	0	0	0	0		
2	Logical Block Address (MSB)							
	[4]							
3	Logical Block Address							
	[4]							
4	Logical Block Address							
	[4]							
5	Logical Block Address (LSB)							
	[4]							
6	0	0	0	0	0	0	0	0
7	Byte Transfer Length (MSB)							
	[5]							
8	Byte Transfer Length (LSB)							
	[5]							
9	0	0	0	0	0	0	Flag	Link [6]

The Read Long command (see Table 5.2.2-9) requests that the target transfers data to the initiator. The data passed during the Read Long command shall include the data bytes, followed by the ECC bytes of the single logical block addressed by the command.

Notes.

- [1] In the CDB, the LUN must be zero.
- [2] A corrected (CORRECT) bit of zero requests that a logical block be read without any ECC correction made by the target. A CORRECT bit of one requests that the data be corrected by ECC, if necessary.

Other error recovery procedures separate from ECC are applied in accordance with the parameter settings given in the Mode Sense Error Recovery Page (page code 01). See also each drive Product Manual for a detailed discussion of drive error recovery procedures.

If the DCR bit of the READ-WRITE Error Recovery Page is equal to one and the CORRECT bit equals one the Read Long Command shall be terminated with CHECK CONDITION status and the sense key shall be set to ILLEGAL REQUEST with an additional sense code of INVALID FIELD IN CDB.

- [3] A relative address (RelAdr) bit of one indicates that the logical block address field is a two's complement displacement. This negative or positive displacement is to be added to the logical block address last accessed on the logical unit to form the logical block address for this command. This feature is only available when linking commands. The feature requires that a previous command in the linked group have accessed a block of data on the logical unit.

A RelAdr bit of zero indicates that the logical block address field specifies the first RelAdr logical block of the range of logical blocks to be operated on by this command.

- [4] The logical block address specifies the logical block at which the read operation shall occur. The most recent data written in the addressed logical block shall be returned.
- [5] The byte transfer length specifies the number of bytes of data that shall be transferred. A transfer length of zero indicates that no bytes shall be transferred. This condition shall not be considered as an error. The byte transfer length requested must be equal to the current block size plus all ECC bytes for a data transfer to occur. If an incorrect number of bytes is stated in the command block, this command terminates with a "Check Condition" status. The correct number of bytes can be determined from the information returned in the extended sense data bytes after issuing the Request Sense command. The Request Sense command shall result in the "Illegal Field In CDB" condition with the "Illegal Request" Sense Key. The extended sense ILI bit shall be set. The extended sense information bytes contain the difference (residue) of the requested length minus the actual length in bytes. (Negative values are indicated by two's complement notation).

- [6] See "Control Byte" paragraph 4.2.6.

5.2.2.10 Write Long Command (3FH)

Table 5.2.2-10. Write Long Command

Bit Byte(s)	7	6	5	4	3	2	1	0
0	0	0	1	1	1	1	1	1
1	Logical Unit No.[1] 0 0 0			0	0	0	0	RelAdr [2]
2	Logical Block Address (MSB)							[3]
3	Logical Block Address							[3]
4	Logical Block Address							[3]
5	Logical Block Address (LSB)							[3]
6	0	0	0	0	0	0	0	0
7	Byte Transfer Length (MSB)							[4]
8	Byte Transfer Length (LSB)							[4]
9	0	0	0	0	0	0	Flag	Link [5]

The Write Long Command (see 5.2.2-10) requests that the target write to the medium the data transferred by the initiator. The data passed during the Write Long Command shall include the data bytes and all ECC bytes to be written to the single logical block addressed in the command. The Read Long command is usually issued before issuing a Write Long command. The Write Long data passed must be in the same order and must be the same number of bytes as the Read Long command.

Notes.

[1] LUN must be zero.

[2] and [3]

A relative address (RelAdr) bit of one indicates that the logical block address field [3] is a two's complement displacement. This negative or positive displacement is to be added to the logical block address last accessed on the logical unit to form the logical block address for this command. This feature is only available when linking commands. The feature requires that a previous command in the linked group have accessed a block of data on the logical unit.

A RelAdr bit of zero indicates that the logical block address field specifies the logical block at which the write operation shall occur.

[4] The byte transfer length specifies the number of bytes of data that shall be transferred. A transfer length of zero indicates that no bytes shall be transferred. This condition shall not be considered as an error. The byte transfer length requested must be equal to the current block size plus all ECC bytes for a data transfer to occur. If an incorrect number of bytes is stated in the command block, this command terminates with a "Check Condition" status. The correct number of bytes can be determined from the information returned in the extended sense data bytes after issuing the Request Sense command. The Request Sense command shall result in the "Illegal Field In CDB" Condition with the "Illegal Request" Sense Key. The extended sense ILI bit shall be set. The extended sense information bytes contain the difference (residue) of the requested length minus the actual length in bytes. (Negative values are indicated by two's complement notation).

[5] See "Control Byte" paragraph 4.2.6.

5.2.2.11 Write Same command (41h)

The Write Same command (Table 5.2.2-11) requests that the target write the single block of data transferred by the initiator to the medium multiple times.

Table 5.2.2-11. Write Same command

Bit Byte(s)	7	6	5	4	3	2	1	0
0	Operation Code (41h)							
1	Logical Unit No. [1]			Reserved		PBdata [2]	LBdata [3]	RelAdr [4]
2	(MSB)							
3	Logical Block Address [5]							
4								
5								
6	Reserved							
7	(MSB)							
8	Number of Blocks [6]							
9	(LSB)							
9	Control [7]							

Notes:

- [1] Logical Unit Number. Always zero.
- [2] A physical block data (PB data) bit of one requests that the target replace the first eight bytes of the data to be written to the current physical sector with the physical address of the sector currently being written using the physical sector format.
- [3] A logical block data (LB data) bit of one requests that the target replace the first four bytes of the data to be written to the current logical block with the logical block address of the block currently being written.
- [4] and [5]
See Notes [3] and [4], Table 5.2.2-9 for a definition of the RelAdr bit and the logical block address field.
- [6] Number of blocks field specifies the number of contiguous logical blocks to be written. A number of blocks field of zero requests that all the remaining logical blocks on the medium be written.
- [7] See "Control Byte" paragraph 4.2.6.

5.2.2.12 Reserve (10) command (56h)

The RESERVE and RELEASE commands provide the basic mechanism for contention resolution in multiple-initiator systems. The RESERVE (10) command (Table 5.2.2-12) is used to reserve a logical unit or, if the extent reservation option is implemented, extents within a logical unit. The third-party reservation allows logical units or extents to be reserved for another specified SCSI device.

Table 5.2.2-12. Reserve (10) Command

Bit Byte(s)	7	6	5	4	3	2	1	0
0	Operation Code (56h)							
1	Logical Unit No. [1] 0 0		0	3rdPty [2]	Reserved			Extent [3]
2	Reservation Identification							[4]
3	Third Party Device ID							[2]
4	Reserved							
5	Reserved							
6	Reserved							
7	(MSB)			Extent List Length			[3]	
8								(LSB)
9	Control							[5]

Notes:

- [1] The Logical Unit Number must be zero.
- [2] If bit 4 is zero, byte 3 must also be zero.
If bit 4 is one, byte 3 identifies the SCSI bus ID of the device for which the drive is reserved.
- [3] Must be zero if not supported by the drive. See drive Product Manual.
- [4] Must be zero if not supported by the drive. See drive Product Manual.
- [5] See "Control Byte" paragraph 4.2.6.

For additional information about the fields of this command see section 5.2.1.8.

5.2.2.13 Release (10) Command (57h)

The RESERVE and RELEASE commands provide the basic mechanism for contention resolution in multiple-initiator systems. The RELEASE command (Table 5.2.2-13) is used to release a previously reserved logical unit, or, if the extent release option is implemented, to release previously reserved extents within a logical unit. It is not an error for an initiator to attempt to release a reservation that is not currently valid. In this case, the drive shall return GOOD status without altering any other reservation.

Table 5.2.2-13. Release (10) Command

Bit Byte(s)	7	6	5	4	3	2	1	0
0	Operation Code (57h)							
1	Logical Unit No. [1] 0 0 0			3rdPty [2]	Reserved			Extent 0 [3]
2	Reservation Identification							[4]
3	Third Party Device ID							[2]
4	Reserved							
5	Reserved							
6	Reserved							
7	Reserved							
8	Reserved							
9	Control Byte							[5]

Notes:

[1] The Logical Unit Number must be ZERO.

[2] If bit 4 is zero, byte 3 is zero. If bit 4 is one, byte 3 identifies the SCSI bus ID of the device that reserves the drive.

[3] Must be zero

[4] Must be zero if not supported. Check with drive Product Manual, "SCSI Interface commands supported" section (see paragraph 5.2.1.8.1 of this volume).

[5] See "Control Byte" paragraph 4.2.6.

For additional explanation about the Release Command, see section 5.2.1.9.

5.2.2.14 Mode Select (10) Command (55h)

The MODE SELECT (10) command provides a means for the initiator to send a list of drive operating mode parameters to the drive. See the MODE SELECT command (section 5.2.1.7) for a description of the fields in this command. Initiators should issue MODE SENSE prior to MODE SELECT to determine supported pages, page lengths, and other parameters.

Table 5.2.2-14. Mode Select (10) Command

Bit Byte(s)	7	6	5	4	3	2	1	0
0	0	1	0	1	0	1	0	1
1	Logical Unit No.			PF [7]	Reserved			SMP [7]
2 : 6	Reserved							
7	(MSB) _____							
8	Parameter List Length							(LSB) _____
9	Control							

Table 5.2.2-14a. Mode Select (10) Parameter List

Bit Byte(s)	7	6	5	4	3	2	1	0
0, 1	Reserved							
0	0	0	0	0	0	0	0	0
2	Medium Type [1]							
	0	0	0	0	0	0	0	0
3,4,5	Reserved							
0	0	0	0	0	0	0	0	0
6,7	Block descriptor length either 0 or 8 (Decimal) [2]							

BLOCK DESCRIPTOR DATA

0	Density Code [3]							
	0	0	0	0	0	0	0	0
1	Number of blocks (MSB) [4]							
	0	0	0	0	0	0	0	0
2	Number of blocks							
	0	0	0	0	0	0	0	0
3	Number of blocks (LSB)							
	0	0	0	0	0	0	0	0
4	0							
5	Block length (MSB)							
6	Block length							[5]
7	Block length (LSB)							

PARAMETER INFORMATION [6]

0 - n	Mode - Select Page Headers and Their Parameters (Tables 5.2.1-14 and 5.2.1-15)
-------	--

Notes: [1] - [6] See Table 5.2.1-13 notes.

[7] See Table 5.2.1-12 notes.

5.2.2.15 Mode Sense (10) Command (5Ah)

The MODE SENSE (10) command provides a means for the drive to report drive operating mode parameters to the initiator. It is a complementary command to the MODE SELECT (10) command. See the MODE SENSE command (Section 5.2.1.10) for a description of the fields in this command.

Table 5.2.2-15. Mode Sense (10) Command

Bit Byte(s)	7	6	5	4	3	2	1	0
0	0	1	0	1	1	0	1	0
1	Logical Unit No.			Rsvd	DBD [8]	Reserved		
2	PCF		Page Code					
3 : 6	Reserved							
7	(MSB)							
8	Allocation Length							
9	(LSB)							
9	Control							

Table 5.2.2-15a. Mode Sense Parameter List

Bit Byte(s)	7	6	5	4	3	2	1	0
0, 1	Sense Data Length [1]							
2	Medium Type							
	0	0	0	0	0	0	0	0 [2]
3	WP[3]	Reserved		DPO- FUA[8]	Reserved			
		0	0		0	0	0	0
4,5	Reserved							
6,7	Block Descriptor Length (8 decimal)							
	0	0	0	0	1	0	0	0 [4]

BLOCK DESCRIPTOR DATA

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

PARAMETER INFORMATION

0-n	Mode Sense Page Headers and Their Parameters							
-----	--	--	--	--	--	--	--	--

See Table 5.2.1-20 for notes [1] thru [7].

See Table 5.2.1-18 [2] for note [8].

5.2.2.16 Pre-Fetch command (34h)

The PRE-FETCH command (Table 5.2.2-16) requests that the drive read and transfer the specified logical blocks to the drive's cache memory. No data shall be transferred to the initiator.

Table 5.2.2-16. Pre-Fetch command

Bit Byte(s)	7	6	5	4	3	2	1	0
0	Operation Code (34h)							
1	Logical Unit No.			Reserved			Immed [1]	RelAdr [2]
2 3 4 5	(MSB) Logical Block Address [3]							(LSB)
6	Reserved							
7 8	(MSB) Transfer Length [4]							(LSB)
9	Control [5]							

Notes:

[1] An immediate (Immed) bit of one indicates that status shall be returned as soon as the command descriptor block has been validated. An Immed bit of zero indicates that status shall be returned after the operation is complete.

[2] and [3]

See 5.2.2.1, notes [2] and [3] for a definition of the RelAdr bit and the Logical Block Address field.

[4] The transfer length field specifies the number of contiguous logical blocks of data that shall be transferred to the drive's cache memory. A transfer length of zero indicates that the contiguous logical blocks up to and including the last logical block of the logical unit shall be transferred to the drive's cache memory. Any other value indicates the number of logical blocks that shall be transferred. The drive may elect to not transfer logical blocks that already are contained in the cache memory.

If the Immed bit is zero and the specified logical blocks were successfully transferred to the cache memory, the drive shall return CONDITION MET status. If the link bit (see 4.2.6) is one, the drive shall return INTERMEDIATE-CONDITION MET status.

If Immed is one, and the unlocked cache memory has sufficient capacity to accept all of the specified logical blocks, the drive shall return CONDITION MET status. If the link bit (see 4.2.6) is one, the drive shall return INTERMEDIATE-CONDITION MET status.

If Immed is one, and the unlocked cache memory does not have sufficient capacity to accept all of the specified logical blocks, the drive shall return GOOD status. The drive shall transfer to cache memory as many logical blocks as will fit. If the link bit (see 4.2.6) is one, the drive shall return INTERMEDIATE status.

[5] See paragraph 4.2.6.

5.3 Xor Commands

The SCSI Xor commands are a special group of commands designed principally for use by array controllers in operating arrays of storage devices, though they are not limited to such a usage. These commands are members of command groups for direct access devices. The Xor commands are grouped here for convenience of reference because of their common functionality, rather than in the command groups where they might otherwise be separately grouped according to their command code numbers. Table 5.3-1 lists the Xor commands described in this section and the subsection number in which each is described. Mode page 10h "Xor control mode page," is used in conjunction with the commands of this section, but is listed in Section 5.2.1 (Table 5.2.1-32) with the other mode page descriptions.

Table 5.3-1. Xor commands defined in this section

Operation Code (hex)	Command name	Subsection	Page
50	XDWRITE	5.3.1	205
51	XPWRITE	5.3.2	206
52	XDREAD	5.3.3	207
80	XDWRITE EXTENDED	5.3.4	208
81	REBUILD	5.3.5	210
82	REGENERATE	5.3.6	213

Not all Seagate devices support the Xor commands. Where used by a particular model Seagate device, the Product Manual for that model device specifies that the Xor commands are supported by the devices addressed by the Product Manual.

In storage arrays, an array controller organizes a group of storage devices into a redundancy group. Some areas within the address space of the storage array are used for check data. The check data is generated by performing a cumulative exclusive-or (xor) operation with the data from other areas within the address space of the storage array known as protected data. This xor operation can be performed by the array controller or by the storage device.

Performing the xor operation in the storage device may result in a reduced number of data transfers across the interconnect. For example, when the xor operation is done within the array controller four data transfer operations are needed for a typical update write sequence: a read transfer from the device containing protected data, a write transfer to the device containing protected data, a read transfer from the device containing check data, and a write transfer to the device containing check data. The array controller also does two internal xor operations in this sequence. In contrast, during array controller supervised xor operations, only three data transfer operations are needed: a write transfer to the device containing protected data, a read transfer from the device containing protected data, and a write transfer to the device containing check data. Note that the array controller doesn't do any internal xor operations. In further contrast, during third party xor operations, only two data transfer operations are needed: a write transfer from the array controller to the device containing protected data and a write transfer from the device containing protected data to the device containing check data. Note that the array controller only issues one command and does no xor operations.

Performing the xor operation in the device eliminates the need for the array controller to perform any xor operations. An array controller performs three basic operations that require xor functionality. These are the update write, regenerate, and rebuild operations.

For additional information on the use of the Xor commands see ANSI document X3T10/94-111r9.

5.3.1 XDWRITE command (50h)

The XDWRITE command (see table 5.3.1-1) requests that the target xor the data transferred to it 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.

Table 5.3.1-1. XDWRITE command

Bit Byte(s)	7	6	5	4	3	2	1	0
0	Operation Code (50h)							
1	Reserved			DPO[1]	FUA[2]	Disable[3] Write	Reserved	
2 3 4 5	(MSB) Logical Block Address [4]							(LSB)
6	Reserved							
7 8	Transfer Length [5]							
9	Control [6]							

The resulting xor data is retained 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 this command.

Notes:

- [1] If the Disable Page Out (DPO) bit is set to one, no data is cached. The DPO bit is only meaningful if the RCD bit of Mode Select Page 8 is set false (Caching enabled).
- [2] A force unit access (FUA) bit of one indicates that the write command shall not return GOOD status until the logical blocks have actually been written on the media. The FUA bit is only meaningful if the WCE bit of Mode Sense page 8 is true.
- [3] A disable write bit of zero indicates that the data transferred from the initiator shall be written to the medium after the xor operation is complete. A disable write bit of one indicates that the data shall not be written to the medium.
- [4] The Logical Block address specifies the logical block at which the write operation shall begin.
- [5] The transfer length field specifies the number of logical blocks that shall be transferred to the XDWRITE target for the Xor operation. It also specifies the number of blocks to be written to the medium after the Xor operation.
- [6] See "CONTROL BYTE," paragraph 4.2.6.

5.3.2 XPWRITE command (51h)

The XPWRITE command (see table 5.3.2-1) requests that the target xor the data transferred to it with the data on the medium and then writes the xor data to the medium.

Table 5.3.2-1. XPWRITE command

Bit Byte(s)	7	6	5	4	3	2	1	0
0	Operation Code (51h)							
1	Reserved			DPO[1]	FUA[2]	Reserved		
2 3 4 5	(MSB) Logical Block Address [3]							(LSB)
6	Reserved							
7 8	Transfer Length [4]							
9	Control [5]							

Notes:

- [1] If the Disable Page Out (DPO) bit is set to one, no data is cached. The DPO bit is only meaningful if the RCD bit of Mode Select Page 8 is set false (Caching enabled).
- [2] A force unit access (FUA) bit of one indicates that the write command shall not return GOOD status until the logical blocks have actually been written on the media. The FUA bit is only meaningful if the WCE bit of Mode Sense page 8 is true.
- [3] The logical block address field specifies the starting logical block address at which the target shall read data from its medium. It also specifies the starting logical block address at which the xor result data is to be written to the medium.
- [4] The transfer length field specifies the number of blocks to be read from the XPWRITE target medium for the Xor operation. It also specifies the number of blocks to be written to the medium after the Xor operation
- [5] See "Control Byte" paragraph 4.2.6.

5.3.3 XDREAD command (52h)

The XDREAD command (see table 5.3.3-1) requests that the target transfer to the initiator the xor data generated by an XDWRITE or REGENERATE command.

Table 5.3.3-1. XDREAD command

Bit Byte(s)	7	6	5	4	3	2	1	0
0	Operation Code (52h)							
1	Reserved							
2 3 4 5	(MSB)	Logical Block Address [1]						(LSB)
6	Reserved							
7 8	Transfer Length [2]							
9	Control [3]							

Notes:

[1][2] The xor data transferred is identified by logical block address and transfer length that are the same as those specified in a prior XDWRITE or REGENERATE 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..

[3] See "Control Byte" paragraph 4.2.6.

5.3.4 XDWRITE EXTENDED command (80h)

The XDWRITE EXTENDED command (see table 5.3.4-1) requests that the target xor the data transferred to it with the data on the medium (see Note [6]). The disposition of the data transferred from the initiator is controlled by the disable write bit. The resulting xor data is sent to a secondary device using an XPWRITE command.

Table 5.3.4-1. XDWRITE EXTENDED command

Bit Byte(s)	7	6	5	4	3	2	1	0
0	Operation Code (80h)							
1	Table[1] Address	Reserved		DPO[2]	FUA[3]	Disable Write[4]	Port control [5]	
2 3 4 5	(MSB)	Logical Block Address [6]						(LSB)
6 7 8 9	(MSB)	Secondary logical block address [7]						(LSB)
10 11 12 13	(MSB)	Transfer length [8]						(LSB)
14	Secondary address [9]							
15	Control [10]							

NOTES:

- [1] A table address bit of zero indicates that the secondary address field contains the target identifier of the [9] target to which the xor data is transferred. The LUN of the secondary target shall be zero.

Note: If the protocol requires more than one byte for the target identifier and the table address bit is set to zero, the secondary address field specifies the least significant byte of the secondary target identifier - the upper bytes of the secondary target identifier are assumed to be equal to the upper bytes of the target identifier of the XDWRITE EXTENDED target.

A table address bit of one indicates that the secondary address field contains a pointer to a look up table of SAM compliant target identifiers. This look up table is reserved for future definition.

- [2] If the Disable Page Out (DPO) bit is set to one, no data is cached. The DPO bit is only meaningful if the RCD bit of Mode Select Page 8 is set false (Caching enabled).
- [3] A force unit access (FUA) bit of one indicates that the write command shall not return GOOD status until the logical blocks have actually been written on the media. The FUA bit is only meaningful if the WCE bit of Mode Sense page 8 is true.

- [4] A disable write bit of zero indicates that the data transferred from the initiator shall be written to the medium after the xor operation is complete. A disable write bit of one indicates that the data shall not be written to the medium.
- [5] The port control field is defined in table 5.3.5-2. If the port control field has a value of 01b and the target is not a multiple port device the command shall be terminated with a CHECK CONDITION status. The sense data shall be set to ILLEGAL REQUEST: INVALID FIELD IN CDB.
- [6] The Logical Block Address field specifies the logical block at which the Read operation of the data to be Xored shall begin.
- [7] The transfer length field specifies the number of logical blocks that shall be transferred to the
- [8] XDWRITE EXTENDED target, and to the XPWRITE target.

The xor data transfer to the secondary target is performed using an XPWRITE command. The XPWRITE command is sent to the device specified in the secondary address field [9]. The secondary logical block address field value [7] is placed in the logical block address field of the XPWRITE command. The transfer length field value is placed in the transfer length field of the XPWRITE command. The completion status of the XDWRITE EXTENDED command shall not be returned to the initiator until the completion status of the XPWRITE command has been received.

Note: The xor data transfer to the secondary target may be broken into multiple XPWRITE commands. If this is done, the XDWRITE EXTENDED target will need to calculate the logical block addresses and transfer lengths for the individual XPWRITE commands. Also, the completion status of the XDWRITE EXTENDED command shall not be returned to the initiator until the completion status of all XPWRITE commands have been received.

- [10] See "Control Byte," paragraph 4.2.6.

If the XPWRITE command terminates with a CHECK CONDITION status and the sense key is not set to RECOVERED ERROR the XDWRITE EXTENDED command shall return CHECK CONDITION status.

5.3.5 REBUILD command (81h)

The REBUILD command (see table 5.3.5-1) requests that the target write to the medium the xor data generated from the specified source devices. The target, acting as a temporary initiator, issues READ commands to retrieve the specified data from the source device.

Table 5.3.5-1. REBUILD command

Bit Byte(s)	7	6	5	4	3	2	1	0
0	Operation Code (81h)							
1	Reserved		DPO[1]	FUA[2]	InData [3]	Port control [4]		
2 3 4 5	(MSB)	Logical Block Address [5]						(LSB)
6 7 8 9	(MSB)	Rebuild length [6]						(LSB)
10 11 12 13	(MSB)	Parameter list length [7]						(LSB)
14	Reserved							
15	Control [8]							

Note: The target that receives the REBUILD command is not one of the source devices. If only one source is specified, then an xor operation does not occur. This case can occur in disk mirroring applications.

If the command terminates with CHECK CONDITION status the sense data shall contain the logical block address of the failed block with the lowest logical block address. All logical blocks affected by the command and having a logical block address lower than that of the reported failing block shall be rebuilt and written to the medium.

Notes:

- [1] If the Disable Page Out (DPO) bit is set to one, no data is cached. The DPO bit is only meaningful if the RCD bit of Mode Select Page 8 is set false (Caching enabled).
- [2] A force unit access (FUA) bit of one indicates that the write command shall not return GOOD status until the logical blocks have actually been written on the media. The FUA bit is only meaningful if the WCE bit of Mode Sense page 8 is true.
- [3] If the intermediate data (IntData) bit is set to zero, then intermediate data is not sent with the rebuild parameter list. If the bit is set to one, the rebuild parameter list includes intermediate data. The length of the intermediate data can be calculated by multiplying the rebuild length times the block size. This data shall be treated as an additional source, and an xor operation performed with it and the data from the specified sources.

- [4] The port control field is defined in table 5.3.5-2. If the port control field has a value of 01b and the target is not a multiple port device the command shall be terminated with a CHECK CONDITION status. The sense data shall be set to ILLEGAL REQUEST: INVALID FIELD IN CDB.

Table 5.3.5-2. Port control field

Value	Description
00	The target transfers the data using the same port that received the command.
01	The target transfers the data using a different port than that which received the command.
10	The target transfers the data using one port of the targets choice.
11	The target transfers the data using one or more ports of the targets choice.

- [5] The logical block address field specifies the starting logical block address at which the target shall write the xor result data on its own medium.
- [6] The rebuild length field specifies the number of blocks to be written to the medium. It also specifies the number of blocks that are read from each source.
- [7] The parameter list length field specifies the length in bytes of the parameter list that shall be transferred from the initiator to the target (see table 5.3.5-3).
- [8] See "Control Byte," paragraph 4.2.5.

The REBUILD parameter data is described in table 5.3.5-3.

Table 5.3.5-3. REBUILD and REGENERATE parameter data

Bit Byte(s)	7	6	5	4	3	2	1	0	
0	Number of source descriptors (x) [1]								
1	Reserved								
2	Source descriptor/pad length (MSB) [2]								
3	Source descriptor/pad length (LSB) [3]								
	Source descriptor(s) (if any) [3]								
4	(LSB)								
19	Source descriptor (first) [3]								
	· · [3] ·								
16x - 12	Source descriptor (last) [3]								
16x + 3									
16x + 4	Pad, if any (Length y) [4]								
16x+y+3									
16x+y+z+4	MSB	Intermediate data, if any (Length z) [5]							
16x+y+z+3									(LSB)

Notes:

- [1] The number of source descriptors field indicates the number of source descriptors in the parameter data.
- [2] The source descriptor/pad length specifies the sum of the lengths in bytes of all of the source descriptors and the pad.
- [3] The source descriptors identify the source device target identifiers and starting logical block addresses on those devices for the regenerate or rebuild operation. See Table 5.3.5-4 for the source descriptor format.
- [4] The pad field contains invalid data and shall be ignored.

Note: The pad field is included to accommodate initiators which require the intermediate data to be aligned on a particular memory boundary.

- [5] The intermediate data field contains data that shall be used in the xor operation with the data from the specified source devices. The length of the data is equal to the rebuild/regenerate length multiplied by the block size.

Table 5.3.5-4. Source descriptor format

Bit Byte(s)	7	6	5	4	3	2	1	0
0 : 7	(MSB) Source device address [1]							(LSB)
8 : 11	(MSB) Reserved							(LSB)
12 : 15	(MSB) Source starting logical block address [2]							(LSB)

Notes:

- [1] The source device address field specifies a SAM compliant target identifier of a device that is a data source.
- [2] The source starting logical block address field indicates the starting logical block address to use when reading data from the source specified in the source device address field.

5.3.6 REGENERATE command (82h)

The REGENERATE command (see Table 5.3.6-1) requests that the target write to the buffer the xor data generated from its own medium and the specified source devices. The target, acting as a temporary initiator, issues READ commands to retrieve the specified data.

Table 5.3.6-1. REGENERATE command

Bit Byte(s)	7	6	5	4	3	2	1	0
0	Operation Code (82h)							
1	Reserved		DPO[1]	FUA[2]	InData [3]	Port control [4]		
2 3 4 5	(MSB) Logical Block Address [5]							(LSB)
6 7 8 9	(MSB) Regenerate length [6]							(LSB)
10 11 12 13	(MSB) Parameter list length [7]							(LSB)
14	Reserved							
15	Control [8]							

The resulting xor data is retained in the target's buffer until it is retrieved by an XDREAD Command with a starting logical block address and transfer length that match the logical block address and regenerate length of this command.

Notes:

- [1] If the Disable Page Out (DPO) bit is set to one, no data is cached. The DPO bit is only meaningful if the RCD bit of Mode Select Page 8 is set false (Caching enabled).
- [2] A force unit access (FUA) bit of one indicates that the write command shall not return GOOD status until the logical blocks have actually been written on the media. The FUA bit is only meaningful if the WCE bit of Mode Sense page 8 is true.
- [3] See table 5.3..5-1 for a definition of the IntData Bit.
- [4] See table 5.3.5-2 for a definition of the port control field.
- [5] The logical block address field specifies the starting logical block address for the target to read data from its own medium. This data is a source for the regenerate operation.
- [6] The regenerate length field indicates the length in logical blocks of the resulting xor data. It also specifies the length in logical blocks that is transferred from each of the specified sources.
- [7] The parameter data for the REGENERATE command is defined in table 5.3.5-3. This parameter data describes the other devices that will be sources for the regenerate operation. The target receiving the REGENERATE command is implicitly a source, and is not included in the parameter data.
- [8] See "Control Byte." paragraph 4.2.6.

6.0 Seagate technical support services

If you need assistance installing your drive, consult your dealer. Dealers are familiar with their unique system configurations and can help you with system conflicts and other technical issues. If you need additional assistance with your Seagate® drive or other Seagate products, use one of the Seagate technical support services listed below.

SeaFONE® 1-800-SEAGATE

Seagate's 800 number (1-800-732-4283) allows toll-free access to automated self-help services, providing answers to commonly asked questions, troubleshooting tips, and specifications for disc drives and tape drives. This service is available 24 hours daily and requires a touch-tone phone. International callers can reach this automated self-help service by dialing 408-456-4496.

Online services

Using a modem, you can obtain troubleshooting tips, free utility programs, drive specifications and jumper settings for Seagate's entire product line. You can also download software for installing and analyzing your drive.

SeaNET™

You can obtain technical information about Seagate products over the Internet from Seagate's World Wide Web home page (<http://www.seagate.com>) or Seagate's ftp server (<ftp://ftp.seagate.com>). You can also send E-mail with your questions to **DiscSupport @ Seagate.com** or **TapeSupport @ Seagate.com**.

Seagate CompuServe forum

Online technical support for Seagate products is available on CompuServe. To access our technical support forum, type **go seagate**. This forum provides information similar to that found on SeaBOARD. In addition, you can type questions or browse through previous questions and answers on the forum messages.

SeaBOARD®

SeaBOARD is a computer bulletin board system that contains information about Seagate disc and tape drive products and is available 24 hours daily. Set your communications software to eight data bits, no parity, and one stop bit (8-N-1).

Location	Phone number
Australia	61-2-9756-2359
England	44-1628-478011
France	33 1-48 25 35 95
Germany	49-89-140-9331
Singapore	TBA
Taiwan	886-2-719-6075
Thailand	662-531-8111
USA	Disc: 408-434-1080; Tape: 408-456-4415

FAX services

SeaFAX®

You can use a touch-tone telephone to access Seagate's automated FAX system to receive technical support information by return FAX. This service is available 24 hours daily.

Location	Phone number
Australia	61-2-9756-5170
England	44-1628-894084
USA	1-800-SEAGATE or 408-456-4496

Seagate technical support FAX

You can FAX questions or comments to technical support specialists 24 hours daily. Responses are sent during business hours.

Location	Phone number
Australia	61-2-9725-4052
England	44-1628-890660
France	33 1-46 04 42 50

Location	Phone number
Germany	49-89-1430-5100
Hong Kong	852-2368 7173
Japan	81-3-5462-2979
Korea	82-2-556-7294/4251
Singapore	65-488-7528
Taiwan	886-2-715-2923
USA	408-944-9120

Direct-support services

Seagate technical support

For one-on-one help, you can talk to a technical support specialist during local business hours. Before calling, note your system configuration and drive model number (STxxxx).

Location	Phone number
Australia	61-2-9725-3366 (9:00 a.m. to 5:00 p.m., M–F)
England	44-1628-894083 (10:00 a.m. to 1:00 p.m., 2:00 p.m. to 5:00 p.m., M–F)
France	33 1-41 86 10 86 (9:30 a.m. to 12:30 p.m., 2:00 p.m. to 5:00 p.m., M–F)
Germany	Disc: 49-89-140-9332; Tape: 49-89-140-9333 (9:30 a.m. to 12:30 p.m., 2:00 p.m. to 4:00 p.m., M–F)
Hong Kong	852-2368 9918
Korea	82-2-556-8241
Singapore	65-488-7584 (9:00 a.m. to 12:00 p.m., 2:00 p.m. to 5:00 p.m., M–F)
Taiwan	886-2-514-2237
USA	Please dial 1-800-SEAGATE or 408-456-4496 for the specific product telephone number. (6:00 a.m. to 11:15 a.m., 12:30 p.m. to 5:00 p.m., Pacific time, M–F)

SeaTDD™ 408-944-9121

Using a telecommunications device for the deaf (TDD), you can send questions or comments 24 hours daily and exchange messages with a technical support specialist between 6:00 a.m. to 11:15 a.m. and 12:30 p.m. to 5:00 p.m. (Pacific time) Monday through Friday.

Customer service centers

Seagate direct OEM, Distribution, and System Integrator customers should contact their Seagate service representative for warranty information. Other customers should contact their place of purchase. Seagate offers comprehensive customer support for all Seagate drives. These services are available worldwide.

Location	Phone number	FAX number
Asia Pacific and Australia	65-485-3595	65-485-4980
Europe, Middle East, and Africa	31-2031-67300	31-2065-34320
Japan	81-3-5462-2904	81-3-5462-2979
USA	1-800-468-3472	405-949-6740
Other Americas (Brazil, Canada, Mexico)	405-949-6706	405-949-6738

Manufacturer's representatives

Brazil		
MA Informatica	55-11-810-7794	55-21-253-6467
Canada		
Memofix	905-660-4936	905-660-8738
Adtech	905-812-8099	905-812-7807
	1-800-624-9857	
Mexico		
Abicom Seamax SA DE CV	525-546-6965	525-546-4888

Appendix A. SCSI Configured Auto Magically (SCAM)

A.1.0 General

SCAM Protocol defines all the hardware and software requirements that provide SCSI devices with a means to establish a device's SCSI bus ID over the SCSI interface using the standard SCSI interface signals in a nontraditional way. The SCAM Protocol defines a level 1 protocol and a level 2 protocol. Level 1 protocol defines a hardware and software functionality that is less capable than Level 2 protocol. Hot plugging is an example of a feature not provided under Level 1 protocol. However, Level 1 still provides a means to solve most configuration problems common to single user systems.

SCAM Protocol possesses a means to isolate each device on the SCSI bus and assign each a separate SCSI bus ID. This is done even though the devices may have no assigned ID (SCAM protocol provides "assigned ID's") or may already have a hard ID established by jumpers or switches. Some of the devices being put on the bus could have the same hard ID. SCAM assigns distinct "soft" ID's so there are no duplicate IDs on the SCSI bus. SCAM tolerance devices are allowed to keep whatever ID ("current ID") they already have (each must be different) and those ID's are not assigned to any other device.

Each SCAM compliant device has a unique identification string of up to 31 bytes assigned at manufacturing time. No two SCSI devices in the world have the same identification string numeric value. The SCAM protocol uses these device identification strings to isolate each device on the SCSI bus one at a time and assign SCSI ID numbers to the different devices on the bus. Thereafter (when the SCAM protocol has been completed) each device on the SCSI bus uses its assigned or current SCSI bus ID to arbitrate for the use of the bus, unless the SCAM protocol becomes operative again for some reason.

Standard Seagate SCSI devices implement SCAM Protocol. As a factory installed option, full SCAM protocol may be turned on so the device is SCAM Level 2 capable. OEM drives as shipped support SCAM Level 1.

For those interested in the details of SCAM theory of operation, a companion Product Manual is available titled "SCSI configured automatically (SCAM)," P/N 77767519. It contains detailed explanations, flow diagrams, tables, timing waveform diagrams and a digital scope readout diagram for a typical SCAM operation.



Seagate Technology, Inc.
920 Disc Drive, Scotts Valley, California 95066-4544, USA
Publication Number: 77738479, Rev. H (08/97)

Printed in USA