



.....
SCSI Interface
.....

.....
Product Manual, Volume 2
.....

.....
SCSI Interface
.....

.....
Product Manual, Volume 2
.....

© 1997-1999 Seagate Technology, Inc. All rights reserved

Publication number: 77738479, Rev. J

April 1999

Seagate, Seagate Technology, and the Seagate logo are registered trademarks of Seagate Technology, Inc. SeaFAX, SeaFONE, SeaBOARD, and SeaTDD 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, and 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 thru 187. Added pages 188 thru 197.
G	PLD:86785	05/02/96	QB	v thru ix, 1 thru 197. Added pages 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.)
J				v thru ix, 1 thru 216. Added pages 217 - 260.

Table of Contents

1.0	Interface requirements	1
1.1	How to use this interface manual	1
1.1.1	Scope of SCSI-3 standards	1
1.1.2	Applicable standards	2
1.2	General interface description	2
1.2.1	Glossary	3
1.2.2	Keywords	8
1.3	Physical interface characteristics	8
1.4	Summary of SCSI commands and messages	9
2.0	SCSI bus	11
2.1	SCSI bus signals	12
2.1.1	Drive select	13
2.1.2	Signal values	13
2.1.3	OR-tied signals	13
2.1.4	Signal sources	13
2.2	SCSI bus timing	13
2.2.1	Arbitration delay (2.4 μ s)	14
2.2.2	Assertion period (90 ns)	14
2.2.3	Bus clear delay (800 ns)	14
2.2.4	Bus free delay (800 ns)	14
2.2.5	Bus set delay (1.8 μ s)	14
2.2.6	Bus settle delay (400 ns)	14
2.2.7	Cable skew delay (10 ns)	14
2.2.8	Data release delay (400 ns)	14
2.2.9	Deskew delay (45 ns)	15
2.2.10	Disconnection delay (200 μ s)	15
2.2.11	Hold time (45 ns)	15
2.2.12	Negation period (90 ns)	15
2.2.13	Reset hold time (25 μ s)	15
2.2.14	Selection abort time (200 μ s)	15
2.2.15	Selection timeout delay (250 ms recommended)	15
2.2.16	Transfer period (negotiated by Synchronous Data Transfer message)	15
2.3	Fast synchronous transfer option	15
2.3.1	Fast assertion period (30 ns)	15
2.3.2	Fast cable skew delay (5 ns)	15
2.3.3	Fast deskew delay (20 ns)	15
2.3.4	Fast hold time (10 ns)	16
2.3.5	Fast negation period (30 ns)	16
3.0	Logical characteristics	17
3.1	SCSI bus phases	17
3.1.1	Bus Free phase	17
3.1.2	Arbitration phase	17
3.1.3	Selection phase	18
3.1.3.1	Nonarbitrating system	18
3.1.3.2	Arbitrating systems	18
3.1.3.3	All systems	18
3.1.3.4	Single initiator option	19
3.1.3.5	Selection time-out procedure	19
3.1.3.6	Arbitration and selection under SCAM protocol	19
3.1.4	Reselection phase	19
3.1.4.1	Reselection procedure	20
3.1.4.2	Reselection timeout procedure	20
3.1.5	Information transfer phases	21
3.1.5.1	Asynchronous information transfer	21

	3.1.5.2	Synchronous data transfer	22
	3.1.6	Command phase	22
	3.1.7	Data phase.	22
	3.1.7.1	Data in phase	22
	3.1.7.2	Data out phase	22
	3.1.8	Status phase	23
	3.1.9	Message phase	23
	3.1.9.1	Message in phase	23
	3.1.9.2	Message out phase	23
	3.1.10	Signal restrictions between phases	23
3.2		SCSI bus asynchronous conditions.	23
	3.2.1	Attention condition	23
	3.2.2	Reset condition	24
3.3		SCSI bus phase sequences	24
	3.3.1	Nonarbitrating system	24
	3.3.2	Arbitrating systems	25
	3.3.3	All systems.	25
3.4		SCSI pointers	25
	3.4.1	Current pointers	26
	3.4.2	Saved pointers.	26
3.5		Message system specification.	27
	3.5.1	General message protocol.	27
	3.5.2	Messages—General	28
	3.5.3	Message details.	29
	3.5.3.1	One and two byte messages	29
	3.5.3.2	Extended Message (01h)	34
3.6		Message exception conditions handling	40
3.7		S.M.A.R.T. system	42
4.0		SCSI commands	43
4.1		Command implementation requirements	43
	4.1.1	Reserved	43
4.2		Command Descriptor Block (CDB)	43
	4.2.1	Operation code	44
	4.2.2	Logical Unit Number (LUN)	45
	4.2.3	Logical block address	45
	4.2.4	Relative address bit	45
	4.2.5	Transfer length.	46
	4.2.6	Control field	46
4.3		Status	47
4.4		Command examples	49
	4.4.1	Single command example	49
	4.4.2	Disconnect example.	50
4.5		Timing examples	51
4.6		Command processing considerations and exception conditions.	66
	4.6.1	Auto Contingent Allegiance or Contingent Allegiance	66
	4.6.1.1	Logical Unit response to Auto Contingent Allegiance or Contingent Allegiance	66
	4.6.1.2	Clearing an Auto Contingent Allegiance condition	67
	4.6.2	Overlapped commands	67
	4.6.3	Incorrect Logical Unit selection	67
	4.6.4	Sense data.	68
	4.6.4.1	Asynchronous Event Reporting	68
	4.6.4.2	Autosense	69
	4.6.5	Unit attention condition	69
	4.6.6	Target Hard Reset	70
	4.6.7	Logical Unit Reset	70
4.7		Queued tasks (formerly “Queued I/O Processes”)	71

4.7.1	Untagged task queuing	71
4.7.2	Tagged task queuing	71
4.8	Parameter rounding	72
4.9	Programmable operating definition	72
4.10	Incorrect initiator connection	73
5.0	Command descriptions	75
5.1	Change Definition command (40h)	78
5.2	Compare command (39h)	79
5.3	Copy command (18h)	79
5.4	Copy and Verify command (3Ah)	79
5.5	Format Unit command (04h)	80
5.5.1	Format Unit parameter definition	81
5.5.2	Format Unit parameter list	83
5.5.2.1	Defect List header	83
5.5.2.2	Initialization Pattern Descriptor	84
5.5.2.3	Defect List formats	85
5.6	Inquiry command (12h)	87
5.6.1	Drive Standard Inquiry Data	88
5.6.2	Vital Product Data pages	91
5.6.3	Unit Serial Number page (80h)	92
5.6.4	Implemented Operating Definition page (81h)	92
5.6.5	Device Identification page (83h)	93
5.6.6	Firmware Numbers page (C0h)	96
5.6.7	Date Code page (C1h)	98
5.6.8	Jumper Settings page (C2h)	99
5.6.9	Device Behavior page (C3h)	101
5.7	Lock-unlock Cache command (36h)	101
5.8	Log Select command (4Ch)	102
5.8.1	Write, Read, and Verify Error Counter pages (02h, 03h, and 05h)	107
5.8.2	Non-Medium Error page (06h)	108
5.8.3	Cache Statistics page (37h)	108
5.8.4	Factory Log page (3Eh)	109
5.9	Log Sense command (4Dh)	110
5.10	Mode Select (6) command (15h)	112
5.10.1	Mode Select parameter list	112
5.10.2	Mode Select page descriptors	114
5.11	Mode Select (10) command (55h)	116
5.11.1	Mode Select (10) parameter list	116
5.12	Mode Sense (6) command (1Ah)	119
5.12.1	Mode Sense page descriptor header	122
5.12.2	Error Recovery page (01h)	123
5.12.3	Disconnect/Reconnect Control page (02h)	126
5.12.4	Format Parameters page (03h)	128
5.12.5	Rigid Drive Geometry Parameters page (04h)	130
5.12.6	Verify Error Recovery page (07h)	132
5.12.7	Caching Parameters page (08h)	134
5.12.8	Control Mode page (0Ah)	137
5.12.9	Notch page (0Ch)	139
5.12.10	Power Condition page (0Dh or 1Ah)	141
5.12.11	Xor Control Mode page (10h)	142
5.12.12	Unit Attention Parameters page (00h)	144
5.12.13	Informational Exceptions Control page (1Ch)	146
5.13	Mode Sense (10) command (5Ah)	149
5.14	Move Medium command (A7h)	152
5.15	Persistent Reserve In command (5Eh)	153
5.15.1	Persistent Reserve In parameter data for Read Keys	154
5.15.2	Persistent Reserve In parameter data for Read Reservations	155

5.16	Persistent Reserve Out command (5Fh)	158
5.16.1	Persistent Reserve Out parameter list	159
5.17	Prefetch command (34h)	161
5.18	Prevent/Allow Medium Removal command (1Eh)	161
5.19	Read command (08h)	162
5.20	Read Buffer command (3Ch)	164
5.20.1	Read Combined Descriptor Header and Data mode (0000b)	165
5.20.2	Read Data mode (0010b)	165
5.20.3	Read Buffer Descriptor mode (0011b)	165
5.20.4	Read Data from Echo Buffer mode (1010b)	166
5.20.5	Echo Buffer Descriptor mode (1011b)	166
5.21	Read Capacity command (25h)	167
5.22	Read Defect Data (10) command (37h)	169
5.23	Read Defect Data (12) command (B7h)	171
5.24	Read Element Status command (B4h)	172
5.25	Read Extended command (28h)	173
5.26	Read Long command (3Eh)	175
5.27	Reassign Blocks command (07h)	176
5.27.1	Reassign Blocks defect list	177
5.28	Rebuild command (81h)	178
5.28.1	Rebuild and Regenerate parameter data	179
5.29	Receive Diagnostic Results command (1Ch)	181
5.29.1	Supported Diagnostic Pages	183
5.29.2	Translate Address page (40h)	183
5.30	Regenerate command (82h)	185
5.31	Release (6) command (17h)	187
5.31.1	Logical Unit Release	187
5.31.2	Extent Release	187
5.31.3	Third Party Release	187
5.32	Release (10) command (57h)	188
5.33	Report LUNs command (A0h)	189
5.34	Request Sense command (03h)	191
5.34.1	Extended Sense Data Format	192
5.34.2	Sense Keys	195
5.34.3	Extended Sense, Additional Sense and Additional Sense Qualifier codes	196
5.34.3.1	Deferred errors	199
5.35	Reserve (6) command (16h)	200
5.35.1	Logical Unit Reservation	200
5.35.2	Extent Reservation	200
5.35.3	Third Party Reservation	201
5.35.4	Superseding Reservations	201
5.35.5	Reservations	201
5.36	Reserve (10) command (56h)	202
5.37	Rezero Unit command (01h)	203
5.38	Search Data Equal command (31h)	203
5.39	Search Data High command (30h)	203
5.40	Search Data Low command (32h)	203
5.41	Seek command (0Bh)	204
5.42	Seek Extended command (2Bh)	205
5.43	Send Diagnostic command (1Dh)	206
5.43.1	Supported Diagnostic page—Send Diagnostic	207
5.43.2	Translate Address page—Send Diagnostic	207
5.44	Set Limits command (33h)	208
5.45	Start/Stop Unit command (1Bh)	209
5.46	Synchronize Cache command (35h)	210
5.47	Test Unit Ready command (00h)	211
5.48	Verify command (2Fh)	212
5.49	Write command (0Ah)	213

5.50	Write and Verify command (2Eh)	214
5.51	Write Buffer command (3Bh)	215
5.51.1	Combined Header and Data mode (0000b)	216
5.51.2	Write Data Only mode (0010b)	216
5.51.3	Download Microcode mode (0100b)	216
5.51.4	Download Microcode and Save mode (0101b)	216
5.51.4.1	Buffer ID and Buffer Offset fields (CDB bytes 2-5)	216
5.51.4.2	Byte Transfer Length (CDB bytes 6, 7, and 8)	217
5.51.5	Download Microcode with Offsets (0110b)	217
5.51.6	Download Microcode with Offsets and Save mode (0111b)	217
5.51.7	Write Data to Echo Buffer mode (1010b)	218
5.52	Write Extended command (2Ah)	219
5.53	Write Long command (3Fh)	221
5.54	Write Same command (41h)	222
5.55	XDRead command (52h)	223
5.55.1	Xor commands	223
5.56	XDWrite command (50h)	225
5.57	XDWrite Extended command (80h)	226
5.58	XPWrite command (51h)	228
6.0	Seagate Technology support services	229
Appendix A.	SCSI Configured Automatically (SCAM)	233

List of Figures

Figure 1.	Functional scope of SCSI-3 standards	1
Figure 2.	SCSI client-server model	3
Figure 3.	SCSI ID bits	11
Figure 4.	Phase sequences with arbitration	25
Figure 5.	Phase sequences without arbitration	25
Figure 6.	Simplified SCSI system	25
Figure 7.	Wide SCSI byte ordering	38
Figure 8.	Single command example	49
Figure 9.	Disconnect example	50

1.0 Interface requirements

1.1 How to use this interface manual

For information about features not included herein that the newer drives may support, refer to the standards listed in Section 1.1.1.

This specification is designed to provide a universal detailed description of the SCSI interface for those drive products whose individual drive's Product Manual, Volume 1, do not contain the details of how the SCSI interface is implemented by that drive.

Note. Individual drive's Product Manual, Volume 1, have tables that specify which SCSI-2/SCSI-3 features the drive implements, what the default parameters are for the various features they implement and which parameters are changeable and which are not.

No attempt is made in this universal specification to specify which descriptions or tables apply to only SCSI-2 or SCSI-3. The combination of this general specification with the details in the individual drive's Product Manual, 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-2/SCSI-3 features that Seagate drives support. Reference must be made back to the individual drive's Product Manual to find out which SCSI-2/SCSI-3 features a particular drive supports.

This specification is Volume 2 of a set of manuals that is made up of an individual drive's Product Manual, Volume 1, and this manual. This Volume 2 manual is referenced by other Volume 1 Product Manuals representing other drives Seagate produces.

1.1.1 Scope of SCSI-3 standards

Figure 1 uses a representative set of specifications to show the functional partitions and the relationships among SCSI-3 standards applicable to drives covered by this product manual.

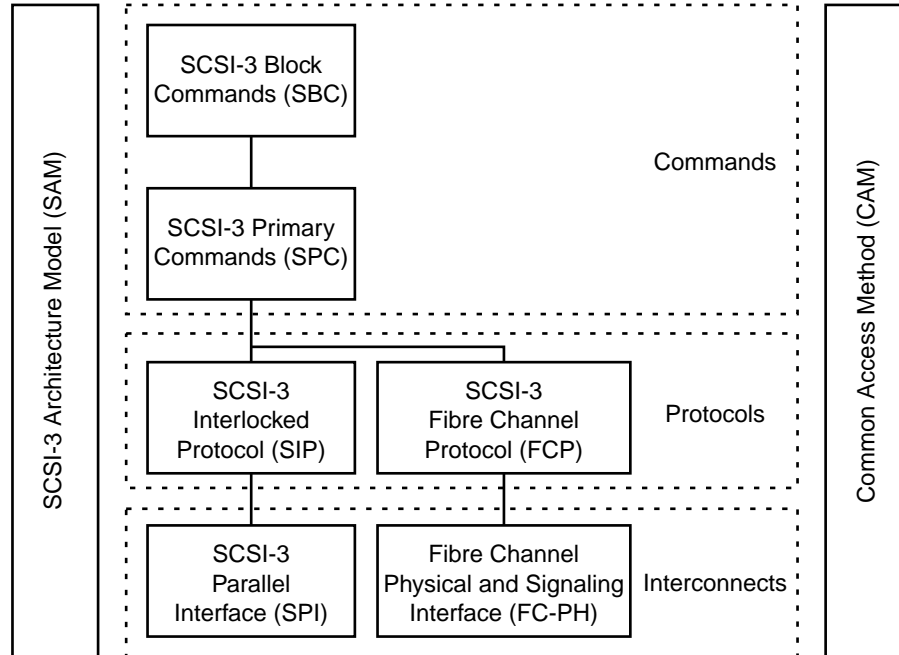


Figure 1. Functional scope of SCSI-3 standards

The functional areas define the scope of each standard as follows:

- **SCSI Architecture Model:** Defines the SCSI systems model, the functional partitioning of the SCSI-3 standard set and requirements applicable to all SCSI-3 implementations and implementation standards.
- **Commands:** Implementation standards which define classes including a device model for each class. These standards specify the required commands and behavior that is common to all devices or unique to a given class of devices and prescribe the rules to be followed by an initiator when sending commands to a device.
- **Common Access Method:** Implementation standard which defines a host architecture and set of services for device access.
- **Protocols:** Implementation standards which define the rules for exchanging information so that different SCSI-3 devices can communicate.
- **Interconnects:** Implementation standards which define the electrical and signaling rules essential for devices to interoperate over a given physical interconnect.

The diagram of Figure 1 shows how the standards listed below fit within each category. The standards included in the diagram are meant to serve as examples and may not reflect the full set of standards currently in force.

1.1.2 Applicable standards

The following ANSI standards should be referenced for more details about SCSI system standards of operation:

- SCSI-3 Architectural Model-2 (SAM-2), T10/1157-D
- Enclosure Services Command Set, NCITS 305-199x, T10/1212-D
- SCSI Primary Commands-2 (SPC-2), T10/1236-D
- SCSI Enhanced Parallel Interface-EPI, T10/1143-D
- SCSI-3 Parallel Interface (SPI-3), T10/1302D
- Profile for Parallel SCSI components used in high availability environments, X3T10/1224-TR
- SCSI-3 Medium Changer Commands (SMC), T10/999-D
- SCSI-3 Controller Command Set-2 (SCC-2), T10/1225-D

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. The interface is compatible with the SCSI Interface Specifications listed in Section 1.1.1. 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’s Product Manual, Volume 1. The bus protocol supports multiple initiators, disconnect/reconnect, additional messages plus 6-byte, 10-byte, and 12-byte Command Descriptor Blocks.

Unless specified otherwise in the individual drive’s Product Manual, Volume 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.

Note. In this revision, some new terminology is introduced as taken from the ANSI specifications. In many instances, the broader scope terms such as “initiator” and “target” are not used, but rather the more specific terms “Application Client” and “Device Server” appear. In Figure 2, it can be seen that several “application clients” from a single initiator may have one or more tasks in queue with several “device servers” in a single target. A drive could be a target or it could be one of the device servers as part of some larger entity. When reading the description, one needs to be able to put the drive of interest in the proper context in terms of what is shown in Figure 2. For a proper understanding of the operation of the SCSI protocol, the terms in the SCSI architectural model as described in ANSI specification T10/

1157-D (SAM-3) should be well understood before reading operation descriptions in any SCSI document. Although a Glossary of terms is provided herein, the definitions may not be adequate for some. The SAM-3 specification gives a more detailed understanding of some of the new SCSI terminology.




Figure 2. SCSI client-server model

1.2.1 Glossary

aborted command—A SCSI command that has been ended by aborting the task created to execute it.

ACA—Auto Contingent Allegiance (see below).

ACA command—A command performed by a task with the ACA attribute (see Section 4.6.1).

application client—An object that is the source of SCSI commands. An object in this sense is not a tangible piece of hardware, but may be a single numeric parameter, such as a logical unit number, or a complex entity that performs a set of operations or services on behalf of another object (see ANSI SAM-2, T10/1157-D).

asynchronous event notification—A procedure used by targets to notify initiators of events that occur when a pending task does not exist for that initiator.

auto contingent allegiance—The condition of a task set following the return of a Check Condition or Command Terminated status.

blocked (task state)—The state of a task that is prevented from completing due to an ACA condition.

blocking boundary—A task set boundary denoting a set of conditions that inhibit tasks outside the boundary from entering the Enabled state.

byte—An 8-bit construct.

call—The act of invoking a procedure.

client-server—A relationship established between a pair of distributed objects where one (the client) requests the other (the server) to perform some operation or unit of work on the client's behalf (see ANSI SAM-2, T10/1157-D).

client—An object that requests a service from a server.

command—A request describing a unit of work to be performed by a device server.

command descriptor block—A structure up to 16 bytes in length used to communicate a command from an application client to a device server.

completed command—A command that has ended by returning a status and service response of Task Complete, Linked Command Complete, or Linked Command Complete (with Flag).

completed task—A task that has ended by returning a status and service response of Task Complete. The actual events comprising the Task Complete response are protocol specific.

confirmation—A response returned to an object, which signals the completion of a service request.

confirmed protocol service—A service available at the protocol service interface, which requires confirmation of completion.

current task—A task that is in the process of sending messages, sending status, transferring data, or transferring command data to or from the initiator.

cyclic redundancy check (CRC)—An error detecting code used to detect the validity of data that has been transferred during the current data phase.

destination device—The SCSI device to which a service delivery transaction is addressed. See source device.

device server—An object within the logical unit which executes SCSI tasks according to the rules for task management described in clause 7 of ANSI SAM-2 document, T10/1157D.

device service request—A request, submitted by an application client, conveying a SCSI command to a device server.

device service response—The response returned to an application client by a device server on completion of a SCSI command.

differential—A signalling alternative that employs differential drivers and receivers to improve signal-to-noise ratios and increase maximum cable lengths.

disconnect—The action that occurs when a SCSI device releases control of the SCSI bus, allowing it to go to the Bus Free phase.

domain—An I/O system consisting of a set of SCSI devices that interact with one another by means of a service delivery subsystem.

dormant (task state)—The state of a task that is prevented from starting execution due to the presence of certain other tasks in the task set.

double transition (DT)—The latching of data on both the assertion edge and the negated edge of the REQ (ACK) signal.

driver—The circuitry used to control the state of the bus.

enabled (task state)—The state of a task that may complete at any time. Alternatively, the state of a task that is waiting to receive the next command in a series of linked commands.

ended command—A command that has completed or aborted.

exception condition—Any event that causes a SCSI device to enter an auto contingent allegiance or contingent allegiance condition.

faulted initiator—The initiator to which a Command Terminated or Check Condition status was returned.

faulted task set—A task set that contained a faulting task.

faulting command—A command that completed with a status of Check Condition or Command Terminated.

faulting task—A task that has completed with a status of Check Condition or Command Terminated.

function complete—A logical unit response indicating that a task management function has finished. The actual events comprising this response are protocol specific.

hard reset—A target response to a reset event or a Target Reset in which the target performs the operations described in Section 4.6.6.

implementation—The physical realization of an object.

implementation-specific—A requirement or feature that is defined in a SCSI-3 standard but whose implementation may be specified by the system integrator or vendor.

implementation option—An option whose actualization within an implementation is at the discretion of the implementor.

initiator—A SCSI device containing application clients which originate device service and task management requests to be processed by a target SCSI device.

interconnect subsystem—One or more physical interconnects which appear as a single path for the transfer of information between SCSI devices in a domain.

intersymbol interference (ISI)—The effect of adjacent symbols on the symbol currently being received.

in transit—Information that has been sent to a remote object but not yet received.

I/O operation—An operation defined by an unlinked SCSI command, a series of linked SCSI commands or a task management function.

I/O process—An I/O process consists of one initial connection or, if information units are enabled, the establishment of a nexus, and a zero or more reconnections, all pertaining to a single task or a group of tasks. An I/O process begins with the establishment of a nexus. An I/O process normally ends with the Bus Free phase if the quick arbitrate method is disabled. If the quick arbitrate method is enabled, the I/O process normally ends with a QA Request message.

I T nexus—A nexus that exists between an initiator and a target.

I T L nexus—A nexus that exists between an initiator, a target, and a logical unit. This relationship replaces the prior I T nexus.

I T L Q nexus—A nexus between an initiator, a target, a logical unit, and a queue tag following the successful receipt of one of the queue tag messages. This relationship replaces the prior I T L nexus.

layer—A subdivision of the architecture constituted by subsystems of the same rank.

linked CDB—A CDB with the link bit in the control byte set to one.

linked command—One in a series of SCSI commands executed by a single task, which collectively make up a discrete I/O operation. In such a series, each command has the same task identifier, and all except the last have the link bit in the CDB control byte set to one.

logical unit—A target-resident entity which implements a device model and executes SCSI commands sent by an application client.

logical unit number—A 64-bit identifier for a logical unit.

logical unit option—An option pertaining to a logical unit, whose actualization is at the discretion of the logical unit implementor.

lower level protocol—A protocol used to carry the information representing upper level protocol transactions.

mandatory—The referenced item is required to claim compliance with a standard.

media information—Information stored within a SCSI device which is non-volatile (retained through a power cycle) and accessible to an initiator through the execution of SCSI commands.

multidrop—A characteristic of the SCSI bus that allows SCSI devices to be connected to the SCSI bus without disrupting the electrical path between the terminators.

multimode single-ended (MSE)—A signalling alternative for multimode SCSI devices that employs MSE drivers and receivers to allow multimode SCSI devices to operate when SE SCSI devices are present on the bus.

nexus—A relationship between an initiator and a target, logical unit, or queue tag that begins with an initial connection and ends with the completion of the associated I/O process. This relationship is formed as the result of a task.

object—An architectural abstraction or “container” that encapsulates data types, services, or other objects that are related in some way.

peer-to-peer protocol service—A service used by an upper level protocol implementation to exchange information with its peer.

peer entities—Entities within the same layer.

pending task—A task that is not a current task.

physical interconnect—A single physical pathway for the transfer of information between SCSI devices in a domain.

port—Synonymous with “service delivery port.” A single attachment to a SCSI bus from a SCSI device.

procedure—An operation that can be invoked through an external calling interface.

protocol—The rules governing the content and exchange of information passed between distributed objects through the service delivery subsystem.

protocol option—An option whose definition within a SCSI-3 protocol standard is discretionary.

protocol service confirmation—A signal from the lower level protocol service layer notifying the upper layer that a protocol service request has completed.

protocol service indication—A signal from the lower level protocol service layer notifying the upper level that a protocol transaction has occurred.

protocol service request—A call to the lower level protocol service layer to begin a protocol service transaction.

protocol service response—A reply from the upper level protocol layer in response to a protocol service indication.

queue—The arrangement of tasks within a task set, usually according to the temporal order in which they were created. *See* task set.

queue tag—The parameter associated with a task that uniquely identifies it from other tagged tasks for a logical unit from the same initiator.

receiver—A client or server that is the recipient of a service delivery transaction.

reference model—A standard model used to specify system requirements in an implementation-independent manner.

request—A transaction invoking a service.

request-response transaction—An interaction between a pair of distributed, cooperating objects, consisting of a request for service submitted to an object followed by a response conveying the result.

request-confirmation transaction—An interaction between a pair of cooperating objects, consisting of a request for service submitted to an object followed by a response for the object confirming request completion.

reset event—A protocol-specific event which may trigger a hard reset response from a SCSI device as described in Section 4.6.6.

response—A transaction conveying the result of a request.

SCSI application layer—The protocols and procedures that implement or invoke SCSI commands and task management functions by using services provided by a SCSI protocol layer.

SCSI device—A device that is connected to a service delivery subsystem and supports a SCSI application protocol.

SCSI device identifier—An address by which a SCSI device is referenced within a domain.

SCSI I/O system—An I/O system, consisting of two or more SCSI devices, a SCSI interconnect and a SCSI protocol, which collectively interact to perform SCSI I/O operations.

SCSI protocol layer—The protocol and services used by a SCSI application layer to transport data representing a SCSI application protocol transaction.

sender—A client or server that originates a service delivery transaction.

server—A SCSI object that performs a service on behalf of a client.

service—Any operation or function performed by a SCSI-3 object, which can be invoked by other SCSI-3 objects.

service delivery failure—Any non-recoverable error causing the corruption or loss of one or more service delivery transactions while in transit.

service delivery port—A device-resident interface used by the application client, device server or task manager to enter and retrieve requests and responses from the service delivery subsystem. Synonymous with “port.”

service delivery subsystem—That part of a SCSI I/O system which transmits service requests to a logical unit or target and returns logical unit or target responses to an initiator.

service delivery transaction—A request or response sent through the service delivery subsystem.

signal—(n) A detectable asynchronous event possibly accompanied by descriptive data and parameters. (v) The act of generating such an event.

source device—The SCSI device from which a service delivery transaction originates. *See* destination device.

subsystem—An element in a hierarchically partitioned system which interacts directly only with elements in the next higher division or the next lower division of that system.

suspended information—Information stored within a logical unit that is not available to any pending tasks.

target—A SCSI device which receives SCSI commands and directs such commands to one or more logical units for execution.

task—An object within the logical unit representing the work associated with a command or group of linked commands. “Work” includes establishing an IT nexus, ITL nexus, ITLQ nexus and all of the other processes associated with the execution of a command by a Logical Unit. Refer to the definitions of “object” in this glossary. *See also* ANSI SAM-2, T10/1157-D, Section 4.9.

task abort event—An event or condition indicating that the task has been aborted by means of a task management function.

task address—An initiator identifies a task to a target using a Task Address. The Task Address object represents either a Tagged Task Address or an Untagged Task Address without regard for the tagged or untagged nature of the Task. A Tagged Task Address is composed of a Logical Unit Identifier and a Tag. An Untagged Task Address is composed of a Logical Unit Identifier.

task completion event—An event or condition indicating that the task has ended with a service response of Task Complete.

task ended event—An event or condition indicating that the task has completed or aborted.

task management function—A task manager service which can be invoked by an application client to affect the execution of one or more tasks.

task management request—A request submitted by an application client, invoking a task management function to be executed by a task manager.

task management response—The response returned to an application client by a task manager on completion of a task management request.

task manager—A server within the target which executes task management functions.

task set—A group of tasks within a target device, whose interaction is dependent on the queuing and auto contingent allegiance rules of Section 4.6.1.

task slot—Resources within the logical unit that may be used to contain a task.

task tags—A Tag is a field containing up to 64 bits that is a component of a Tagged Task Identifier. An initiator assigns tag values in each Tagged Task Identifier in a way that ensures that the identifier uniqueness requirements stated in ANSI SAM-2, T10/1157-D, Section 4.9 are met.

third-party command—A SCSI command which requires a logical unit within the target device to assume the initiator role and send a SCSI command to a target device.

transaction—A cooperative interaction between two objects, involving the exchange of information or the execution of some service by one object on behalf of the other.

unconfirmed protocol service—A service available at the protocol service interface, which does not result in a completion confirmation.

unlinked command—A SCSI-3 command having the link bit set to zero in the CDB control byte.

upper level protocol—An application-specific protocol executed through services provided by a lower level protocol.

1.2.2 Keywords

Several keywords are used to differentiate between different levels of requirements and optionality, as follows:

vendor-specific—Specification of the referenced item is determined by the device vendor.

protocol-specific—Implementation of the referenced item is defined by a SCSI-3 protocol standard (see Section 1.1.1.)

expected—A keyword used to describe the behavior of the models specified by this standard.

mandatory—A keyword indicating items required to be implemented as defined by this standard.

may—A keyword that indicates flexibility of choice with no implied preference.

obsolete—A keyword indicating items that were defined in prior SCSI standards but have been removed from this standard.

option, optional—Keywords that describe features which are not required to be implemented by this standard. However, if any optional feature defined by the standard is implemented, it shall be implemented as defined by the standard.

reserved—A key word referring to bits, bytes, words, fields, and code values that are set aside for future standardization. Their use and interpretation may be specified by the future extensions to this or other standards. A reserved bit, byte, word, or field shall be set to zero, or in accordance with a future extension to this standard. The recipient shall not check reserved bits, bytes, words, or fields. Receipt of reserved code values in defined fields shall be treated as an error.

shall—A keyword indicating a mandatory requirement. Designers are required to implement all such mandatory requirements to ensure interoperability with other standard conformant products.

should—A keyword indicating flexibility of choice with a strongly preferred alternative. Equivalent to the phrase “it is recommended.”

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 drive’s Product Manual, Volume 1, 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 5.0.

Command name	Hex code	Device type	Page number
Change Definition	40	All	78
Format Unit	04	dir. access	80
Inquiry	12	All	87
Log Select	4C	All	102
Log Sense	4D	All	110
Mode Select (6 byte)	15	dir. access	112
Mode Select (10 byte)	55	dir. access	116
Mode Sense (6 byte)	1A	dir. access	119
Mode Sense (10 byte)	5A	dir. access	149
Persistent Reserve In	5E	dir.access	154
Persistent Reserve Out	5F	dir. access	159
Prefetch	34	dir. access	161
Read	08	dir. access	162
Read Buffer	3C	All	164
Read Capacity	25	dir. access	167
Read Defect Data (10 byte)	37	dir. access	169
Read Defect Data (12 byte)	B7	dir. access	171
Read Extended	28	dir. access	173
Read Long	3E	dir. access	175
Reassign Blocks	07	dir. access	176
Rebuild	81	dir. access	178
Receive Diagnostic Results	1C	All	181
Regenerate	82	dir. access	185
Release (6 byte)	17	dir. access	187
Release (10 byte)	57	dir. access	188
Report LUNs	A0	dir. access	189
Request Sense	03	All	191
Reserve (6 byte)	16	dir. access	200
Reserve (10 byte)	56	dir. access	202
Rezero Unit	01	dir. access	203
Seek	0B	dir. access	204
Seek Extended	2B	dir. access	205
Send Diagnostic	1D	All	206
Start/Stop Unit	1B	dir. access	209
Synchronize Cache	35	dir. access	210
Test Unit Ready	00	All	211
Verify	2F	dir. access	212
Write	0A	dir. access	213
Write and Verify	2E	dir. access	214
Write Buffer	3B	All	215
Write Extended	2A	dir. access	219
Write Long	3F	dir. access	221
Write Same	41	All	222
XD Read	52	dir. access	223
XD Write	50	dir. access	225
XD Write Extended	80	dir. access	226
XP Write	51	dir. access	228

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	29
Abort Tag	0D	29
Bus Device Reset	0C	29
Clear Queue	0E	30
Command Complete	00	30
Continue I/O Process	12	30
Disconnect	04	30
Extended Message	01	34
Identify	80-FF	31
Ignore Wide Residue	23	39
Initiate Recovery	0F	not supported
Initiator Detected Error	05	31
Linked Command Complete	0A	33
Linked Command Complete (with flag)	0B	33
Message Parity Error	09	32
Message Reject	07	32
Modify Data Pointer	01	32, 34 (extended message)
No Operation	08	32
Queue Tag Messages		32
Head of Queue Tag	21	33
Ordered Queue Tag	22	33
Simple Queue Tag	20	33
Release Recovery	10	not supported
Restore Pointers	03	33
Save Data Pointers	02	33
Synchronous Data Transfer Request	01	33, 34 (extended message)
Target Transfer Disable	13	34
Wide Data Transfer Request	01	37 (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’s Product Manual, Volume 1, in the “Interface requirements” section, 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 3. 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’s Product Manual, Volume 1, “Option/configuration headers” section). Drives that support SCAM protocol (SCSI Configured Automatically) can have their drive ID assigned via interface also (see Appendix A).

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’s Product Manual, Volume 1.



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

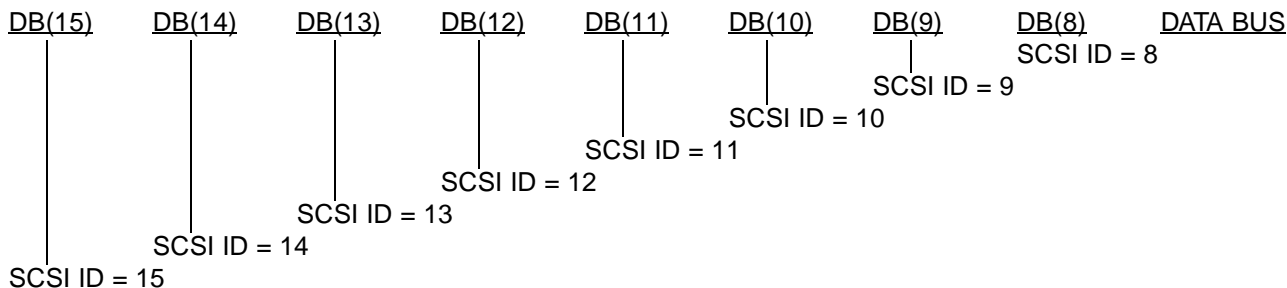


Figure 3. SCSI ID bits

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 3.1.5.2 and 3.5.3.2.

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 • DIFFSNS (Multimode) (may sometimes be designated "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 individual drive's Product Manual, Volume 1, 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.

DIFFSNS (Differential Sense)—When the drive has high voltage differential SCSI I/O circuits (HVD), the DIFFSNS 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.

Multimode—SE or LVD alternative—"LW" and "LC" models have I/O circuits that can operate either in single-ended (SE) or low voltage differential mode (LVD). When the interface "DIFFSNS" line is between 0 V and 0.6 V, the drive interface circuits operate single-ended. When "DIFFSNS" is between +0.7 V and +1.9 V, the drive interface circuits operate low voltage differential. This arrangement is not intended to allow dynamically changing transmission modes, but rather to prevent incompatible devices from attempting to interoperate. Drives must operate only in the mode for which the installation and interface cabling is designed. Multimode I/O circuits used by "LC" and "LW" devices do not operate at high voltage differential levels and should never be exposed to high voltage differential environments unless the command mode voltages in the environment are controlled to safe levels for single-ended and low voltage differential devices (see the ANSI SPI-3 specification, T10/1302D).

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 individual drive's Product Manual, Volume 1. 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 the individual drive's Product Manual, Volume 1. Refer to the "Physical Interface" section of the individual drive's 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. Drives that support SCAM protocol (SCSI Configured Automatically) can have their drive ID assigned via the SCSI interface (see Appendix A).

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 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 2.2.16, 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 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
Information transfer phases					
Command	Target	None	Target	Init.	Init.
Data In	Target	None	Target	Init.	Target
Data Out	Target	None	Target	Init.	Init.
Status	Target	None	Target	Init.	Target
Message In	Target	None	Target	Init.	Target
Message Out	Target	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 ≥ 200 ns use timing parameters specified in Section 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 3.0 are supported by the drive, as specified in each individual drive's Product Manual, Volume 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 Section 3.5.3.1).
5. after a Command Complete message is successfully transmitted from a target (see Section 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.

See also paragraph 3.1.3.6 for arbitration in a SCAM compliant system.

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 Section 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 Section 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 Section 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 time-out 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 time-out procedure is specified for clearing the SCSI bus. If the initiator waits a minimum of a selection time-out 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.

3.1.3.6 Arbitration and selection under SCAM protocol

SCAM selection—SCSI bus signals SEL and MSG asserted, with BSY and DB7–DB0 negated defines the SCAM selection state. This state merely informs the SCAM devices on the SCSI bus of the start of SCAM protocol. The SCAM selection state comes after an arbitration period wherein a SCAM device wins arbitration. Any Level 2 SCAM device can start the SCAM arbitration/selection states. Selection state does not actually select any device as the regular SCSI selection process does, since no device puts its SCSI bus ID on the data lines. See SCAM initialization description in the SCAM Product Manual, part number 77767519. Upon recognizing SCAM selection, a SCAM device's SCSI interface should respond by asserting SEL and MSG itself, then interrupting the device's processor. SCAM target devices shall recognize and respond to SCAM selection whenever the device has neither been assigned a soft ID nor confirmed its current ID since power-on or a reset condition.

SCAM initiator—A SCAM device that is capable of initiating SCAM selection and performing the normal functions of a SCSI initiator. These capabilities permit a SCAM initiator to scan a SCSI bus to discriminate between SCAM tolerant and SCAM devices and assign IDs to the SCAM devices.

SCAM target—A SCAM device that is capable of recognizing and responding to SCAM selection. This capability permits a SCAM target to receive an ID assignment from a SCAM initiator. A SCAM target shall have a current ID, even when a SCAM initiator has not assigned an ID to the SCAM target.

SCAM tolerant—A SCSI device that does not implement the SCAM protocol but complies with certain requirements specified by the SCAM Product Manual. SCAM initiators can detect SCAM tolerant devices on the SCSI bus. SCAM tolerant devices may be intermixed with SCAM devices on the SCSI bus. SCAM tolerant (also called "legacy") devices retain their own established SCSI bus ID.

Unassigned ID—The current SCSI ID (see above definition of "current SCSI ID") that is available to the device (hard set on the drive), but has not yet been assigned by SCAM protocol to the drive.

SCAM requirements—Refer to Appendix A and to the SCAM Product Manual for requirements for devices participating on a SCSI bus that supports the SCAM protocol.

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 Disconnect/Reconnect privilege is granted.

ATN not asserted during Selection means that the host doesn't support the Identify message, so the Disconnect/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 2). 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 2: Information transfer phases

Signal			Phase name	Direction of transfer	Comment
MSG	C/D	I/O			
0	0	0	Data out	Initiator to target	Data Phase
0	0	1	Data In	Initiator from target	
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 Phase
1	1	1	Message In	Initiator from target	

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.

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

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 in Section 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.

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.

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 within a single message phase. Messages supported by a particular drive are listed in the individual drive's Product Manual, Volume 1, 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 Section 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, in Section 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 Section 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 (see Section 3.2.1) and Reset condition (see Section 3.2.2).

3.2 SCSI bus asynchronous conditions

The SCSI bus has two asynchronous conditions: the Attention condition and the Reset 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 in Section 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.

3.3 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

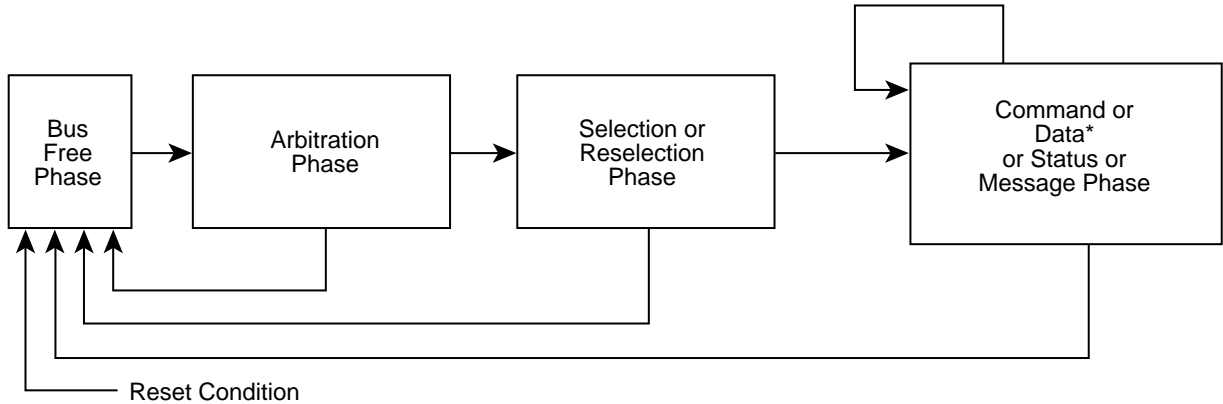
For systems in which the Arbitration phase is not implemented, the allowable sequences are shown in Figure 5. The normal progression is from the Bus Free phase to Selection, and from Selection to one or more of the information transfer phases (Command, Data, Status, or Message).

3.3.2 Arbitrating systems

For systems in which the Arbitration phase is implemented, the allowable sequences are shown in Figure 4. 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). See ANSI document T10/1302D, section 11.1.2.2.1, for a description of Quick Arbitration (QA) phase.

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).



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

Figure 4. Phase sequences with arbitration

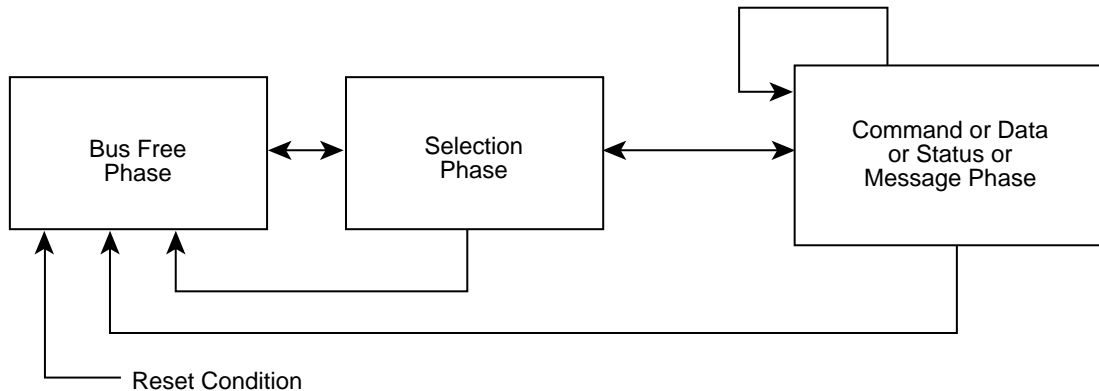


Figure 5. Phase sequences without arbitration

3.4 SCSI pointers

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

Consider the system shown in Figure 6 in which an initiator and target communicate on the SCSI bus in order to execute a command.




Figure 6. Simplified SCSI system

The SCSI architecture provides for sets of pointers to be kept within each initiator Path Control area (see Figure 6). 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 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 (see note following), 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.

Note. 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 Seagate drives, the Logical Unit Number (LUN) (see Section 4.2.2) has always been zero. This may not be true on all future drives.

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. One-byte, Two-byte, and extended message formats are defined. The first byte of the message determines the format as follows:

Message code	Message format
00h	One-Byte Message (Task Complete)
01h	Extended Messages
02h - 1 Fh	One-Byte Message
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 Message code of the byte determines the message to be performed as defined in Table 3. The Identify message is a one byte code, but its format is different.

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. 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 individual drive's Product Manual, Volume 1. Table 3 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: Message codes

Code	Message name	Direction		Negate ATN before last ACK
		In	Out	
01h***	Extended Message	In	Out	Yes
06h	Abort Task Set		Out	Yes
0Dh	Abort Task		Out	Yes
0Ch	Bus Device Reset		Out	Yes
0Eh	Clear Task Set		Out	Yes
00h	Task Complete	In		—
04h	Disconnect	In		—
80h–FFh [2]	Identify	In		—
80h–FFh [2]	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 Operations		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
12h	Continue Task		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 Section 3.2.1, Attention Condition)
 — = Not applicable
 *** = Extended message (see Tables 6 through 9)
 80h+ = Codes 80h through FFh are used for Identify messages
 [1] = The drive does not resend (Retry) the original message
 [2] = See Table 4.

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 Task Set (06h)

The Abort Task Set message is sent from the initiator to the target to clear the present task plus any queued task 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 task sets 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 task and the task set 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 task or task set. 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 Task Set, and Abort Task messages provide a means to clear one or more tasks prior to normal termination. The Bus Device Reset message clears all tasks for all initiators on all logical units of the target. The Clear Queue message clears all tasks for all initiators on the specified logical unit of the target. The Abort Task Set message clears all tasks for the selecting initiator on the specified logical unit of the target. The Abort Task message clears the current task 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 Task (0Dh)

The target shall go to the Bus Free phase following successful receipt of the Abort Task message. The target shall clear the current task for the I T L nexus. If the target has already started execution of the task, 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 task shall be cleared and no status or ending message shall be sent to the initiator. Pending status, data, and tasks for other queued tasks or executing tasks shall not be affected. Execution of other tasks 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 tasks 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 Section 4.6.5).

Clear Task Set (0Eh)

The target shall go to the Bus Free phase following successful receipt of the Clear Task Set message. The target shall perform an action equivalent to receiving a series of Abort Task Set messages from each initiator. All tasks from all initiators in the task set for the specified logical unit shall be cleared from the task set. All executing tasks 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 tasks. A unit attention condition shall be generated for all other initiators with tasks that either had been executing or were queued in a task set for execution for that logical unit. When reporting the Unit Attention condition the additional sense code shall be set to Task Set Cleared by Another Initiator.

Task Complete (00h)

The Task 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 Task Complete message with the ATN signal false.

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

Continue Task (12h)

The Continue Task message is sent from the initiator to the target to reconnect to a task. 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, Task Attribute (if any), and Continue Task messages.

The purpose of the Continue Task message is to distinguish a valid initiator reconnection from an incorrect initiator reconnection (see Section 4.10).

If the target expects a significant delay before it will be ready to continue processing the reconnected task, 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 task for the nexus) and the target shall go to the Bus Free phase (see unexpected disconnect in Section 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 TrnDis bit in the standard Inquiry data (see Table 49).

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 task. 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 Point–Disconnect message sequence.

Identify (80h - FFh)

The Identify message (Table 4) is sent by either the initiator or the target to establish an I T L nexus.

Table 4. Identify message format

Bit Byte	7	6	5	4	3	2	1	0
0	Identify	DiscPriv	LUNTAR	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 (DIImm) 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 task, 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 Unexpected Disconnect, in Section 3.1.1).

Only one logical unit number or target routine number shall be identified per task. 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 task. See the DTDC field of the Disconnect-Reconnect Control Page (Table 91) 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.

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 Section 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, Extended Message.

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.

Task Attribute messages (0Ah, 0Bh, 20h, 21h, 22h or 24h)

When one or more initiators have multiple I/O processes (called tasks) to be queued by a target, each task must have a task tag associated with it. The task tag is specified in a Task Attribute message that must accompany the initiation of each such task. See also Section 4.7.2, "Task Management."

Table 5. Task Attribute message format

Bit Byte	7	6	5	4	3	2	1	0
0	Message Code							
1	Tag (00h-FFh)							

Table 5 defines the format for the queue tag messages. Those drives that implement tagged task queuing and use the Task Attribute messages Head of Queue Task Tag, Ordered Task Set Tag, and Simple Task Tag indicate so in the individual drive's Product Manual, Volume 1.

The Task Attribute messages are used to specify an identifier, called a task tag, for a task that establishes the I L Q nexus. The task tag field is an 8-bit unsigned integer assigned by the initiator during an initial connection. The task tag for every task for each I T L nexus should be unique. If the drive receives a task tag identical to one that is currently in use for the I T L nexus, the drive shall abort all tasks for the initiator and shall return

Check Condition status. The sense key shall be set to Aborted Task and the additional sense code shall be set to Overlapped Tasks Attempted. Only one status is returned.

Note. For each logical unit on each target, each initiator has up to 256 task tags to assign to tasks. A task tag becomes available for reassignment when the task ends. Typically drives to which this specification applies queue up to 64 tasks concurrently for up to seven initiators on the bus. There could be 64 task tags for one initiator, or the 64 can be divided among the seven possible initiators, not necessarily evenly. See individual drive's Product Manual, Volume 1, for number of task tags supported.

The appropriate task 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 task. Only one I T L Q nexus may be established during a connection. If a Task Attribute message is not sent, then only an I T L nexus is established for the task (untagged command).

If a target attempts to reconnect using an invalid task tag, then the initiator should respond with an Abort Task message.

- **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.

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

The Head of Queue task attribute message specifies that the task be placed first in that logical unit's task set for execution. A task already being executed by the target shall not be preempted. A subsequent task received with a Head of Queue message shall be placed at the head of the task set queue for execution in last-in, first-out order.

- **Ordered Task (22h)**

The Ordered Task message specifies that the task be placed in the drive's task queue for execution in the order received, with respect to other commands with Ordered Task messages, except for tasks received with a Head of Queue message, which are placed at the head of the task set queue.

- **Simple Task (20h)**

The Simple Task message specifies that the task be placed in the drive's task set queue for execution. The order of execution can be arranged by the drive in accordance with a performance optimization algorithm. The Simple Task Tag message is also sent by the target when reconnecting to the initiator.

- **ACA Message (24h)**

Auto Contingent Allegiance

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 task) 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 task.

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 task. (See Section 3.4 for a definition of pointers.)

Synchronous Data Transfer Request (01h)

See Paragraph 3.5.3.2 on Extended Message.

Target Transfer Disable (13h)

The Target Transfer Disable (TTD) message is sent from an initiator to a target to request that subsequent reconnections for data transfer on the task 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 Task 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 task, 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 task.

When the target is ready to transfer data for a disconnected task for which a TTD message has been sent, the target shall reconnect to the initiator for the task (via a Reselection phase, an Identify message, and an optional Simple Task 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 task 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 task as described in the Continue Task message (see page 30) 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 TrnDis bit in the standard Inquiry data (see Paragraph 5.6).

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 Pointer message

The Modify Data Pointer message (Table 6) 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.

Table 6: Modify Data Pointer message

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)

Synchronous Data Transfer Request message

A synchronous data transfer agreement only applies to the two SCSI devices that negotiate the agreement. Separate synchronous data transfer agreements are negotiated for each pair of SCSI devices. The synchronous data transfer agreement only applies to data phases.

A Synchronous Data Transfer Request (SDTR) message (Table 7) 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,
3. after a power cycle, and
4. for SCSI devices that implement WDTR, occurrence of an intervening WDTR.

Table 7: Synchronous Data Transfer Request message format

Byte	Value	Description
0	01h	Extended message
1	03h	Extended message length
2	01h	Synchronous Data Transfer Request code
3	[3]	Transfer period factor [1]
4	[2]	REQ/ACK offset

Notes.

- [1] See tables in the individual drive's Product Manual, Volume 1, synchronous data transfer periods supported section, for a list of transfer periods supported by the drive described therein.
- [2] The value supported by a particular drive model is listed in the individual drive's Product Manual, Volume 1, REQ/ACK offset section.
- [3] Transfer factors for this field are listed in Table 8.

Table 8: Transfer Period Factor field

Code	Description
00h-09h	Reserved [1]
0Ah	Transfer period equals 25 ns [2]
0Bh	Transfer period equals 30,3 ns [2]
0Ch	Transfer period equals 50 ns [3]
0Dh-18h	Transfer period equals the transfer period factor * 4 [3]
19h-31h	Transfer period equals the transfer period factor * 4 [4]
32h-FFh	Transfer period equals the transfer period factor * 4 [5]

Notes.

- [1] Faster timings may be allowed by future SCSI parallel interface standards.
- [2] Fast-40 data transfer rates that have a period = 25 ns or 30,3 ns.
- [3] Fast-20 data transfer rates that have a period of ≤ 96 ns and ≥ 50 ns.
- [4] Fast-10 data transfer rates that have a period of ≤ 196 ns and ≥ 100 ns.
- [5] Fast-5 data transfer rates that have a period of ≤ 1020 ns and ≥ 200 ns.

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. This agreement only applies to Data In and Data Out phases. All other phases shall use asynchronous transfers.

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 assertions allowed to be outstanding before the leading edge of its corresponding ACK assertion 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 and that the transfer period factor shall be ignored. A value of FFh shall indicate unlimited REQ/ACK offset. The REQ/ACK offset value supported by each drive is listed in each individual drive's Product Manual, Volume 1, REQ/ACK offset section.

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, smaller 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 response	Implied agreement
1. Non-zero REQ/ACK offset	Synchronous transfer (i.e., 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. Parity error (on responding message)	Asynchronous transfer
5. Unexpected bus free (as a result of the responding message)	Asynchronous transfer
6. 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 Section 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 Section 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 individual drive's Product Manual, Volume 1, REQ/ACK offsets section.

Wide Data Transfer Request message

A Wide Data Transfer Request (WDTR) message (Table 9) exchange shall be initiated by a 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.

Table 9. Wide Data Transfer Request message

Bit Byte	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							

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

In addition, a SCSI device may initiate a 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 a 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 a 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 a 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 Mes-

sage 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 Section 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 Section 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 in Figure 7.




Figure 7. 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 message (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 10. Ignore Wide Residue message format

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

The Ignore Wide Residue message (Table 10) 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 11.

Note. More than one Ignore Wide Residue message may occur during an I/O process.

Table 11: 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

Table 12 attempts to clarify how to handle message exception conditions. The numbers in the grid spaces of Table 12 are response code numbers that are decoded and explained in note [4].

Table 12: SCSI message handling

Message	Phase [1]	Sel	ID	Mout	Cmd	M-in	Resel	Disc	Data	Stat	Cplt
	Response Code [4]										
Abort (06)	2	2	2	2	2	2	2	2	2	2	2
Abort Tag (0D)	4	2	2	2	2	2	2	2	2	2	2
Bus Device Reset (0C)	2	2	2	2	2	2	2	2	2	2	2
Bus Device Reset—Other Ports (14) [2]	4	1	1	1	1	1	1	9,1	1	1	9,1
Clear Task Set (0E)	4	2	2	2	2	2	2	2	2	2	2
Continue Task (12)	4	1	1	3,1	3,1	3,1	3,1	3,9,1	3,1	3,1	3,9,1
Initiator Detected Error (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	C
Identify (Invalid - LUNTAR bit or a reserved bit is set)	3,4	3,1	3,1	3,1	3,1	3,1	3,1	3,9,1	3,1	3,1	3,9,1
Identify (Valid) [3]	1	1	1	1	1	1	1	9,1	1	1	9,1
Message Parity Error (09)	4	4	C	C	5	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	1	9,1	1	1	9,1
Task Set—Simple (20)	4	1	3,1	3,1	3,1	3,1	3,1	3,9,1	3,1	3,1	3,9,1
Task Set—Ordered (22)	4	1	3,1	3,1	3,1	3,1	3,1	3,9,1	3,1	3,1	3,9,1
Task Set—Head (21)	1,4	1	3,1	3,1	3,1	3,1	3,1	3,9,1	3,1	3,1	3,9,1
Synchronous Transfer Request	4	1	1	1	1	1	1	9,1	1	1	9,1
Target Transfer Disable (13)	4	1	1	3,1	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	1	9,1	1	1	9,1
Invalid, unimplemented, illegal messages	4	3,1	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	

[1] Phase abbreviations are as follows:

Abbreviation	Phase/message
Sel	Selection
ID	Identify (after Selection)
Mout	Message Out
Cmd	Command
M-in	Message In (not Cplt or Disc)
Resel	Reselect (after Identify)
Disc	Message In (Disconnect)
Data	Data In/Out
Stat	Status
Cplt	Message In (Task Complete)

[2] 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.

[3] 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.

[4] Responses by number follow:

1	Continue (see Note [3] for multi-initiator or queued environment considerations). The Continue response indicates that the target will continue execution of the current task. 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 task.
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. The Unexpected Bus Free response includes going directly to bus free (ATN signal is ignored if present) and terminating the current task. Since no sense data will be set up, a Contingent Allegiance condition will not be established.
5	<p>Retry Message phase. 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.</p> <p>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.</p>
6	Send Restore Pointers message, and retry phase.
7	Go to Status phase and report Check Condition status (see Note [3] for multi-initiator or queued environment considerations). Associated sense data will have a sense key of Aborted Task (0Bh). The sense code will either indicate an Initiator Detected Error message was received or that a Parity Error was detected, as applicable.
8	Continue with no disconnect (see Note [3] for multi-initiator or queued environment considerations). This is a special case of response 1. The target should continue the I/O process without releasing the bus. The target may attempt to disconnect at a later time.
9	Resend message (Disconnect or Task Complete). 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.
A	If Message In is Save Data Pointer, do response 8. Otherwise, do response 1.
B	Do not respond to the selection.
C	<p>Unexpected Bus Free with sense data set up. 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 Task (0Bh) and a sense code of Invalid Message Error (49h). A Contingent Allegiance condition will be established.</p> <p>Exception: 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.</p>
N/A	Not 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 Sections 5.10, 5.11, 5.12, and 5.13 for the description of the Mode Select/Mode Sense commands and Section 5.12.13 for more details on the Informational Exceptions Control page.

Refer to the individual drive's Product Manual, Volume 1, 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.

Future implementations of the SCSI protocol will contain an autosense feature (see Paragraph 4.6.4.2) allowing the application client to request the automatic return of sense data. Fibre Channel protocol already has this feature in its packetized command structure.

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 13, 14, and 15.

4.2.1 Operation code

The operation code (Table 13) 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 5.0.

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

- Group 0—Six byte commands (see Table 14)
- Group 1—Ten byte commands (see Table 15)
- Group 2—Ten byte commands (see Table 15)
- Group 3—Reserved
- Group 4—Reserved
- Group 5—Twelve byte commands
- Group 6—Vendor specific
- Group 7—Vendor specific

Table 13. Operation code format for CDB

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

Table 14. Typical Command Descriptor Block for six byte commands

Bit Byte	7	6	5	4	3	2	1	0
0	Operation Code							
1	Logical Unit No.			Logical Block Address (if required) (MSB)				
	0	0	0					
2	Logical Block Address (if required)							
3	Logical Block Address (if required)							
4	Transfer Length (if required)							
5	Control Byte							

Table 15. Typical Command Descriptor Block for ten byte commands

Bit Byte	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 (pre-J revision) has been zero. However, future drives may support a non-zero LUN value.

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 and 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 5.21.

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 field

The Control field is the last byte of every command descriptor block. The Control field is defined in Table 16.

Table 16: Control field

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

[1] The Normal Auto Contingent Allegiance (NACA) bit is used to control the rules for handling an Auto Contingent Allegiance (ACA) condition caused by the command. Section 4.6.1.1 specifies the actions to be taken by the logical unit in response to an ACA condition for NACA bit values of one or zero. All drives implement support for the NACA value of zero and may support the NACA value of one. The ability to support an NACA value of one is indicated in standard Inquiry data. See Table 49.

If the NACA bit is set to a value that is not supported, the drive completes the command with a status of Check Condition and a sense key of Illegal Request. The procedure for handling the resulting ACA condition operates in accordance with the supported bit value.

[2] The Link bit allows the drive to continue the task across multiple commands. The Flag bit may be used in conjunction with the Link bit to notify the initiator in an expedited manner that the command has completed.

Support for the Link bit is a logical unit option. A Link bit of one indicates that the initiator requests continuation of the task across two or more SCSI commands. If the Link bit is one and the Flag bit is zero and if the command completes successfully, a drive that supports the Link bit continues the task and returns a status of Intermediate or Intermediate–Condition Met and a service response of Linked Command Complete (see Section 4.3).

Support for the Flag bit is a logical unit option. If the Link bit and Flag bit are both set to one, and if the command completes with a status of Intermediate or Intermediate–Condition Met, a drive that supports the Flag bit returns a service response of Linked Command Complete (with Flag).

The drive completes the command with a status of Check Condition and a sense key of Illegal Request if:

- The Link bit is set to one and the drive does not support linked commands, or
- The Flag bit is set to one and the drive does not support the Flag bit, or
- The Flag bit is set to one and the Link bit is set to zero.

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 17 and 18 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 Section 3.1.1)
5. an Abort Task message
6. a Clear Task Set message

Table 17: Status byte

Bit Byte	7	6	5	4	3	2	1	0
0	Reserved		Status Byte Code					Reserved

Table 18: Status byte code bit values

Bits of status byte								Status represented
7	6	5	4	3	2	1	0	
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/Task Set Full [1]
R	R	1	1	0	0	0	R	ACA Active
All other codes								Reserved

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

[1] What was formerly called a “Command Queue” is now called a “Task Set.”

A description of the status byte codes is given below.

Good—This status indicates that the target has successfully completed the task.

Check Condition—Any error, exception, or abnormal condition that causes sense data to be set shall cause a Check Condition status. An Auto Contingent Allegiance condition has occurred and the Request Sense command should be issued following a Check Condition status to determine the nature of the condition.

Condition Met—This status shall be returned whenever the requested operation specified by an unlinked command is satisfied (see the Search Data and Prefetch commands). This status does not break a chain of linked commands. The logical block address of the logical block that satisfies a 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 (i.e., no reservation conflicts). 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 an Auto

Contingent Allegiance condition exists. The normal initiator recovery action is to issue the command again at a later time.

Intermediate or Intermediate–Condition Met—This status shall 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 Task Set Full or Busy status to be sent. If this status is not returned, the chain of linked commands is broken; no further commands in the series are executed.

Reservation Conflict—The status shall be returned whenever an initiator attempts to access a logical unit, an extent within a logical unit or an element of a logical unit that is reserved with a conflicting reservation type for another SCSI initiator (see the Reserve, Release, Persistent Reserve Out and Persistent Reserve In commands). The recommended initiator recovery action is to issue the command again at a later time. Removing a persistent reservation belonging to a failing initiator may require the executive of a Persistent Reserve Out command with the Preempt or Preempt and Clear actions.

Command Terminated—This status shall be returned whenever the logical unit terminates a task in response to a Terminate Task task management request (see Task Management Functions). This status also indicates that an auto contingent allegiance has occurred (see Task Management Functions).

Task Set Full (formerly “Queue Full”)—This status shall be implemented if the logical unit supports the creation of tagged tasks. This status shall be returned when the logical unit receives a command and does not have enough resources to enter the associated task in the task set.

ACA Active—This status shall be returned when an auto contingent allegiance exists within a task set and an initiator issues a command for that task set when at least one of the following is true:

- a. There is a task with the ACA attribute in the task set;
- b. The initiator issuing the command did not cause the ACA condition;
- c. The task created to execute the command did not have the ACA attribute and the NACA bit was set to one in the CDB control byte of the faulting command (see Task Management Functions).

The initiator may reissue the command after the ACA condition has been cleared.

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 8.

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.

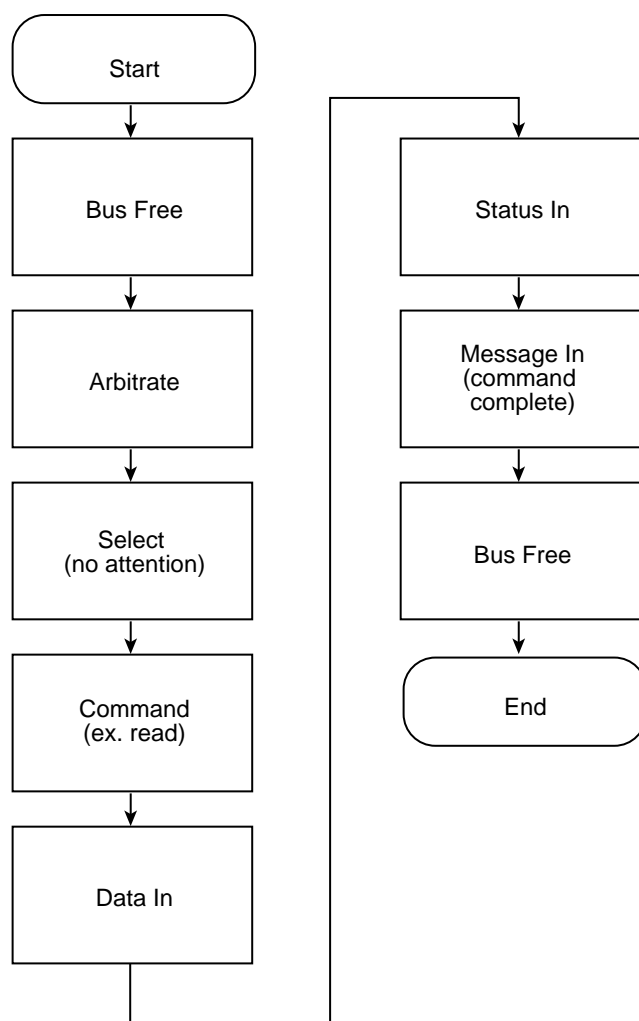


Figure 8. Single command example

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 9.

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.

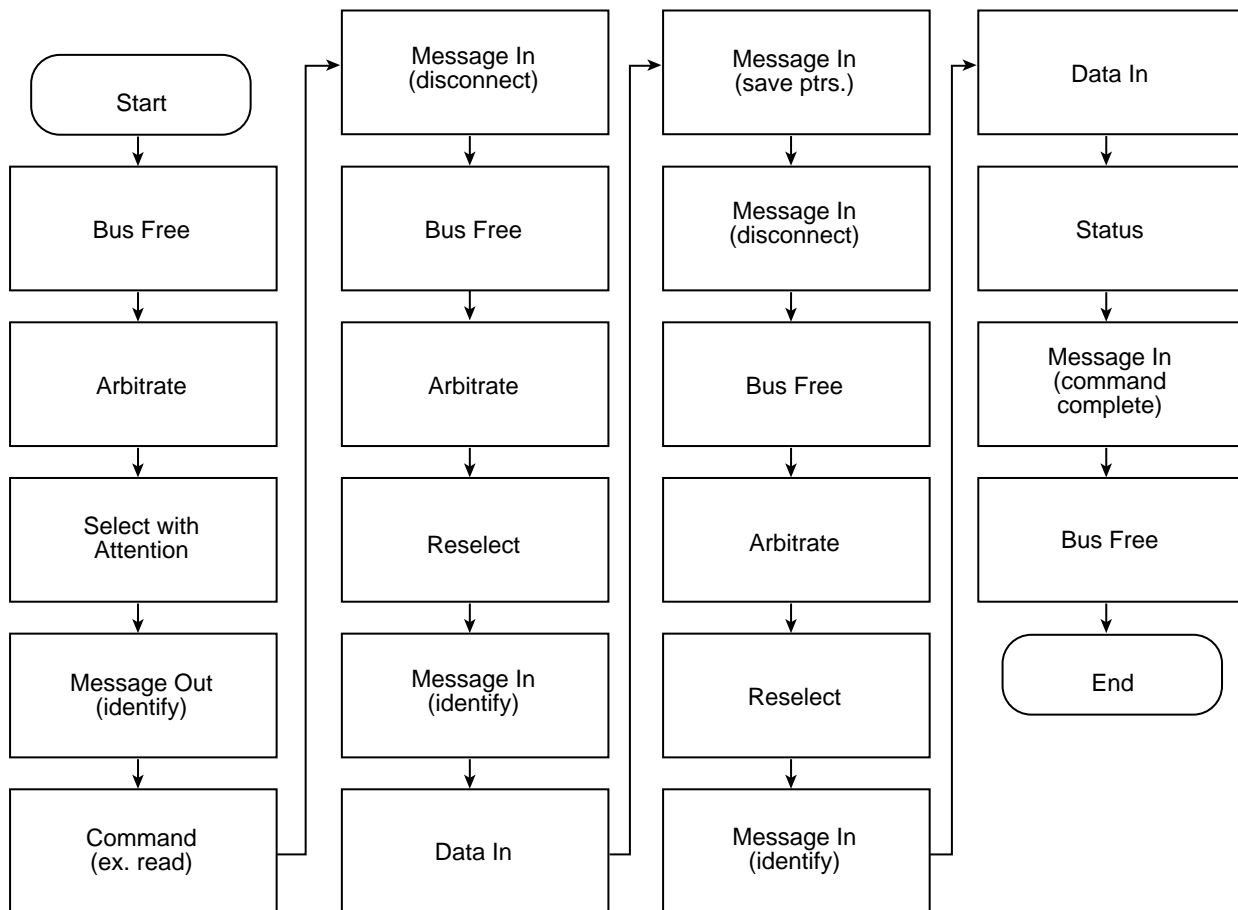



Figure 9. Disconnect example

4.5 Timing examples

Times (T00 through T35) necessary to define performance are listed in the individual drive's Product Manual, Volume 1, in the "Disc drive SCSI timing" section. Timing waveforms to define these times are illustrated in Tables 19 through 33.

Table 19: Arbitration, Selection (No ATN), and Command Phase



* and 8-15 if applicable

Table 20: Arbitration, Selection (with ATN), and Message Out


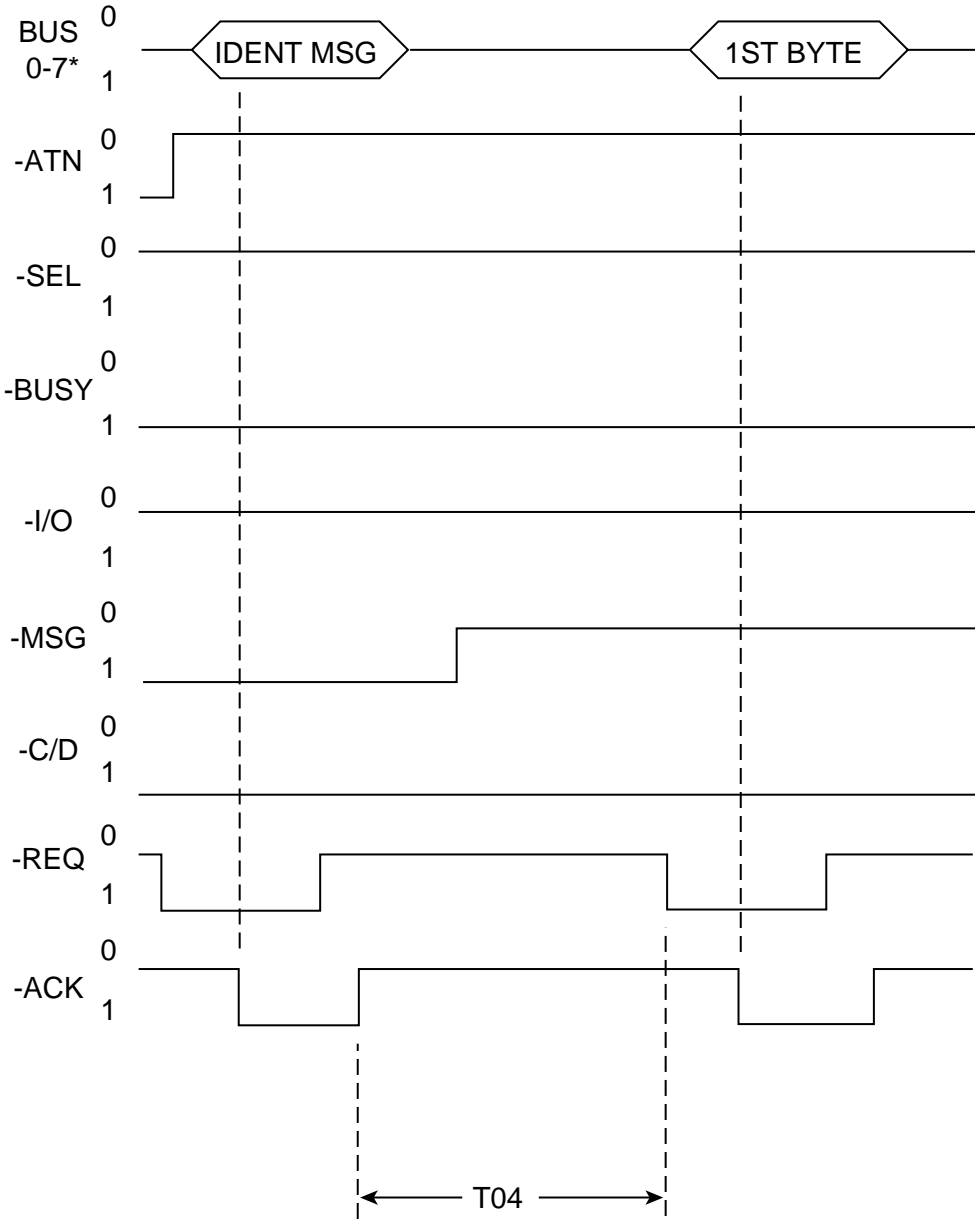
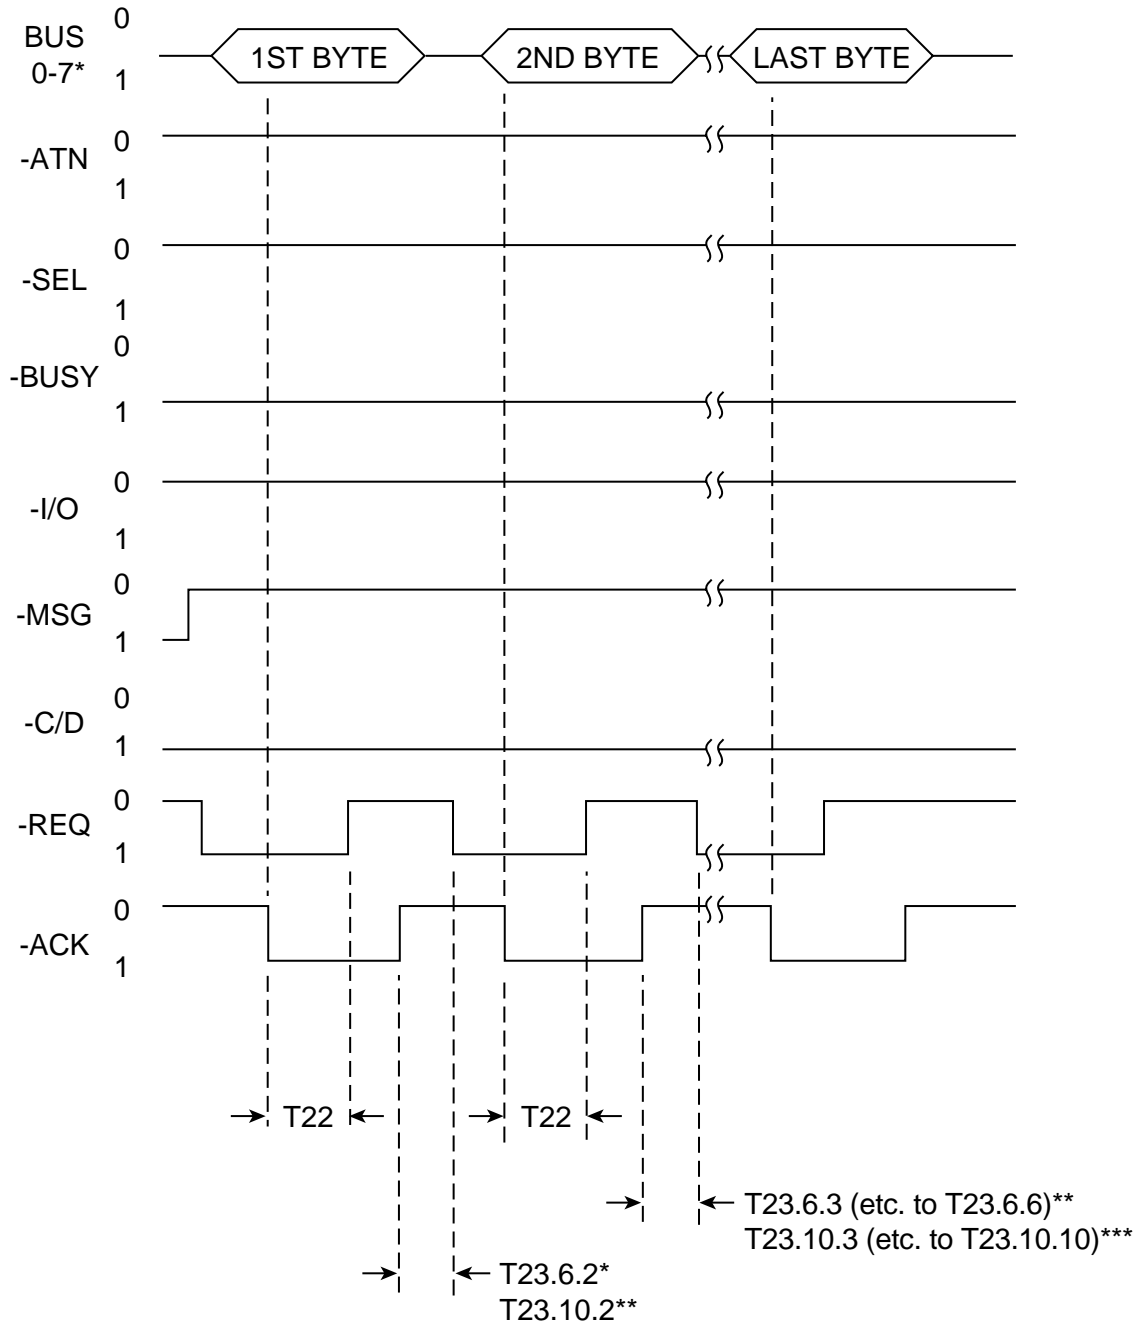


Table 21: Identify Msg Out to Command Phase



* and 8-15 if applicable

Table 22: Command Descriptor Block Transfer




* and 8-15 if applicable

** six byte commands

*** ten byte commands

Table 23: Command Phase, Status Phase, Command Completed Msg and Bus Free



* and 8-15 if applicable

Table 24: Last Command Byte, Disconnect Msg, Bus Free, and Reselect

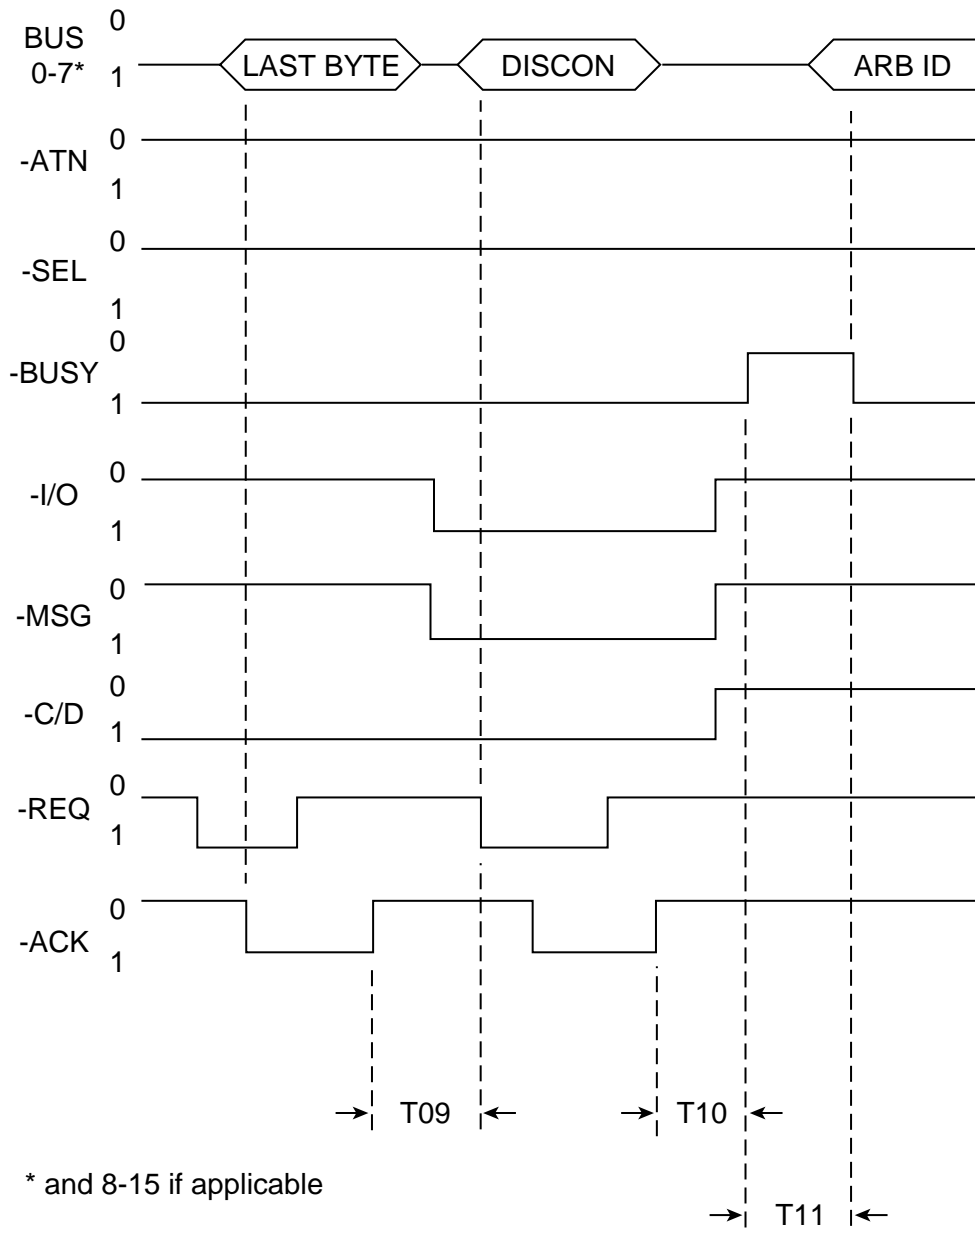




Table 25: Arbitration, Reselection, and Message In



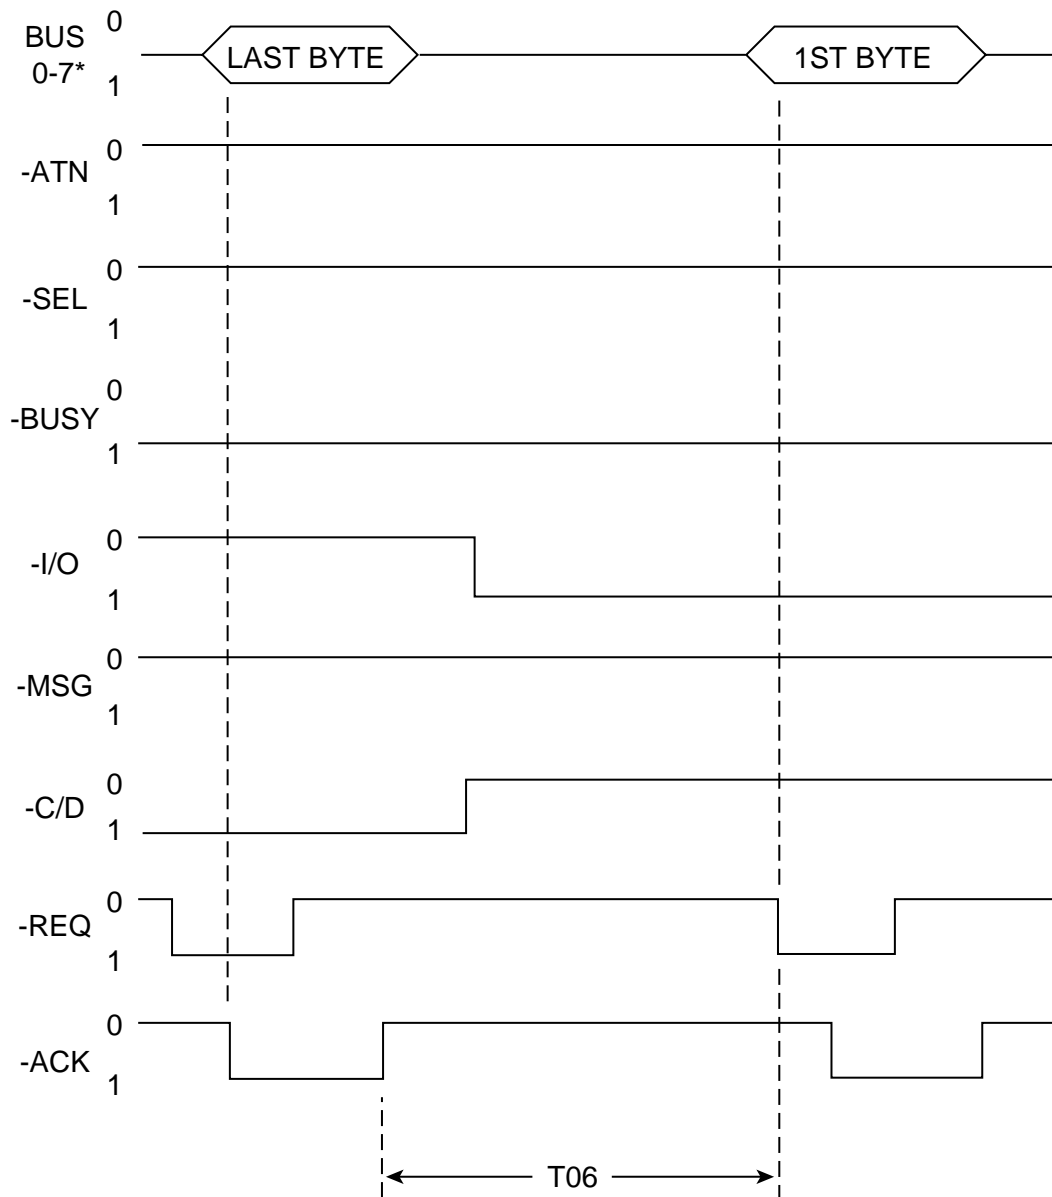
* and 8-15 if applicable

Table 26: Reselect Identify Msg, Status Phase, Command Complete Msg, and Bus Free




* and 8-15 if applicable

Table 27: Last Command Byte to Data in Phase




* and 8-15 if applicable

Table 28: Last Command Byte to Data Out Phase



* and 8-15 if applicable

Table 29: Reselect Identify Msg to Data in Phase



* and 8-15 if applicable

Table 30: Data in Block Transfer

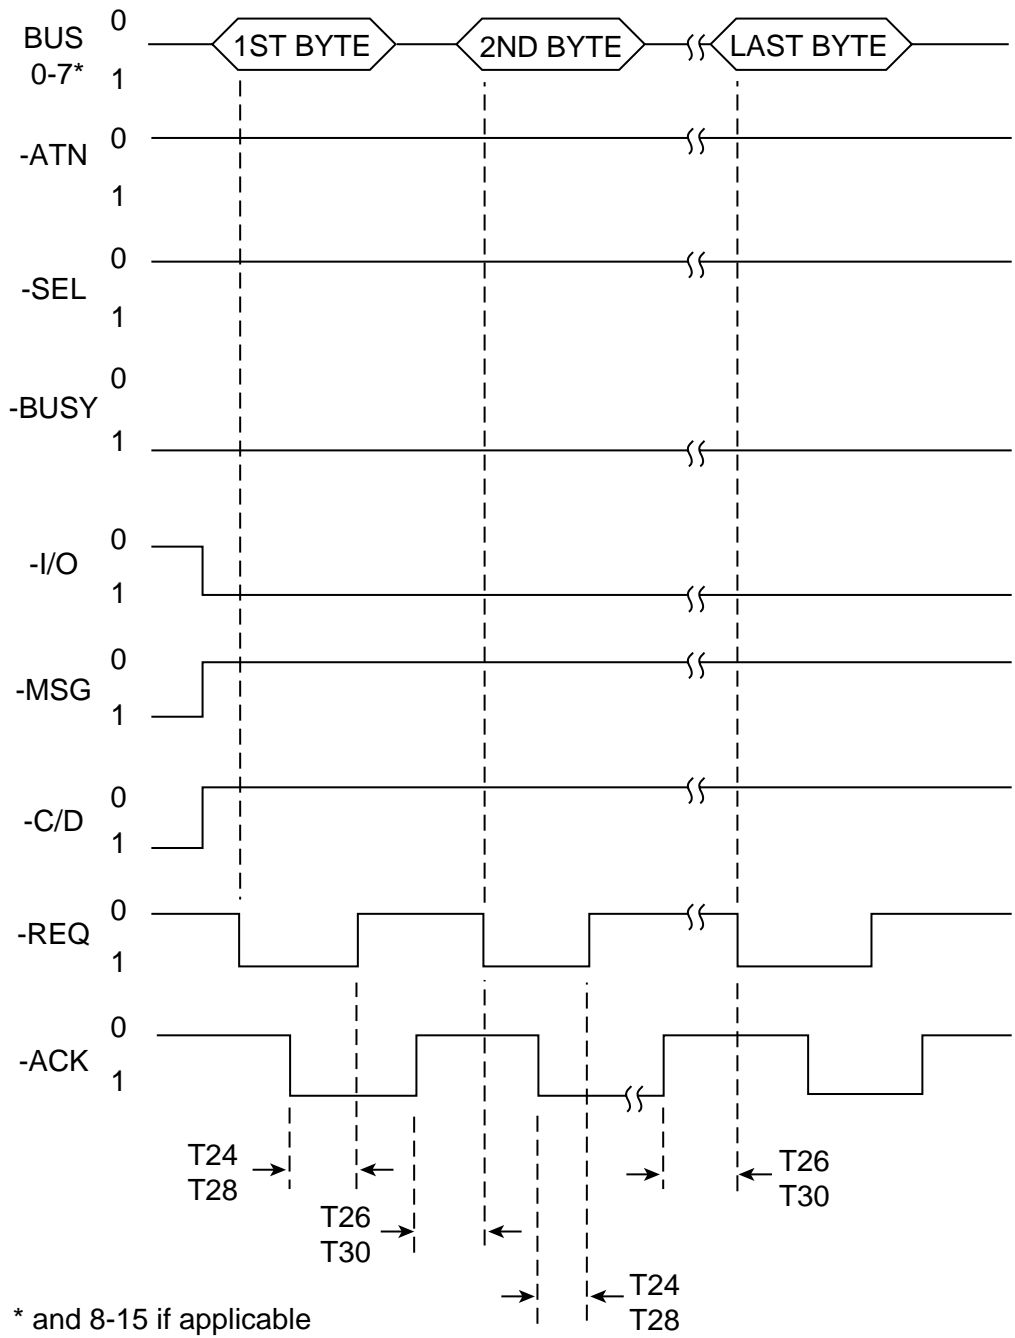
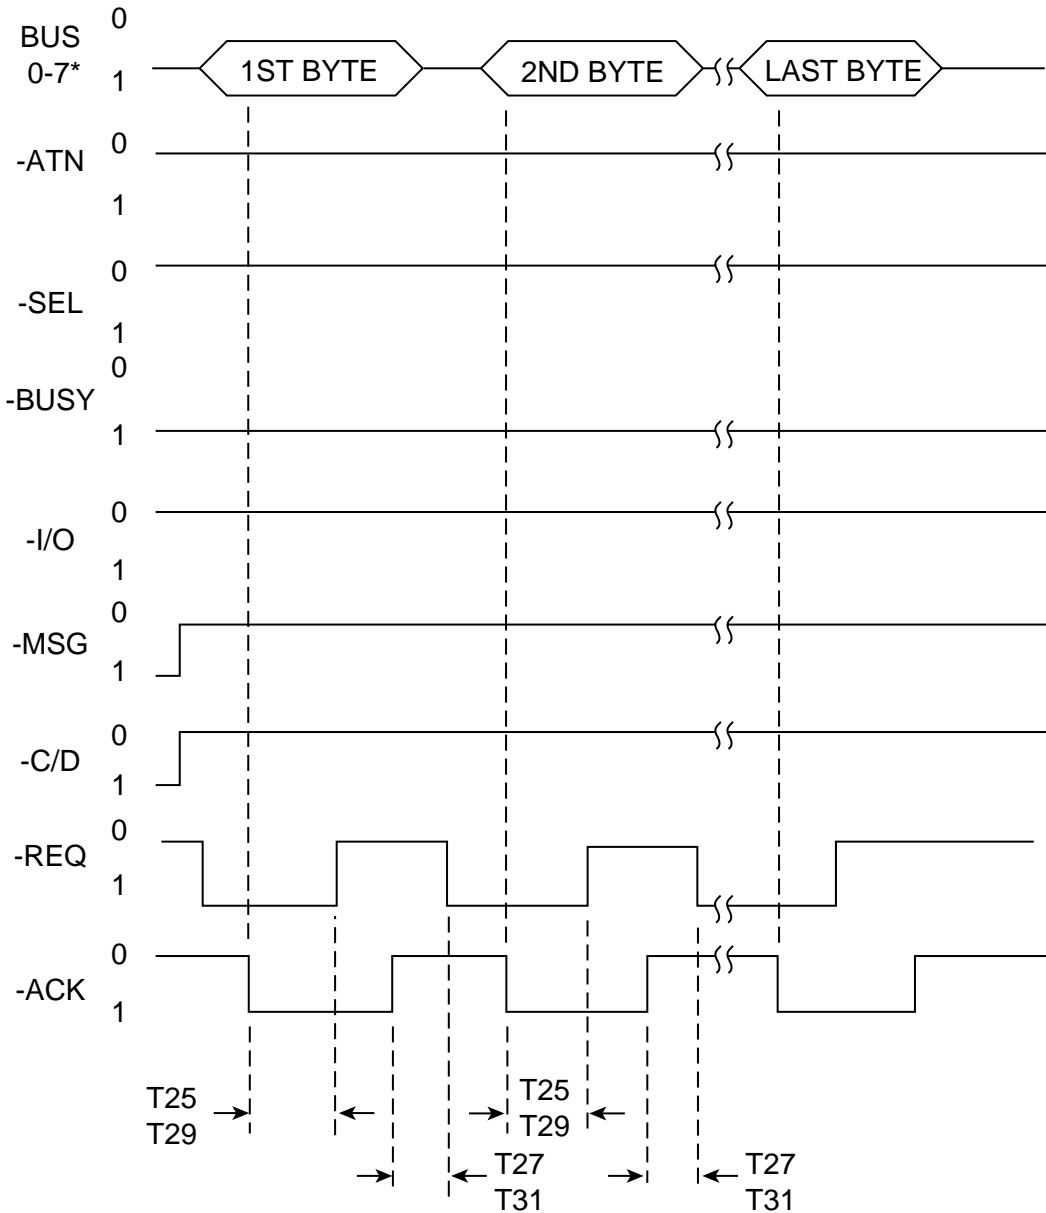
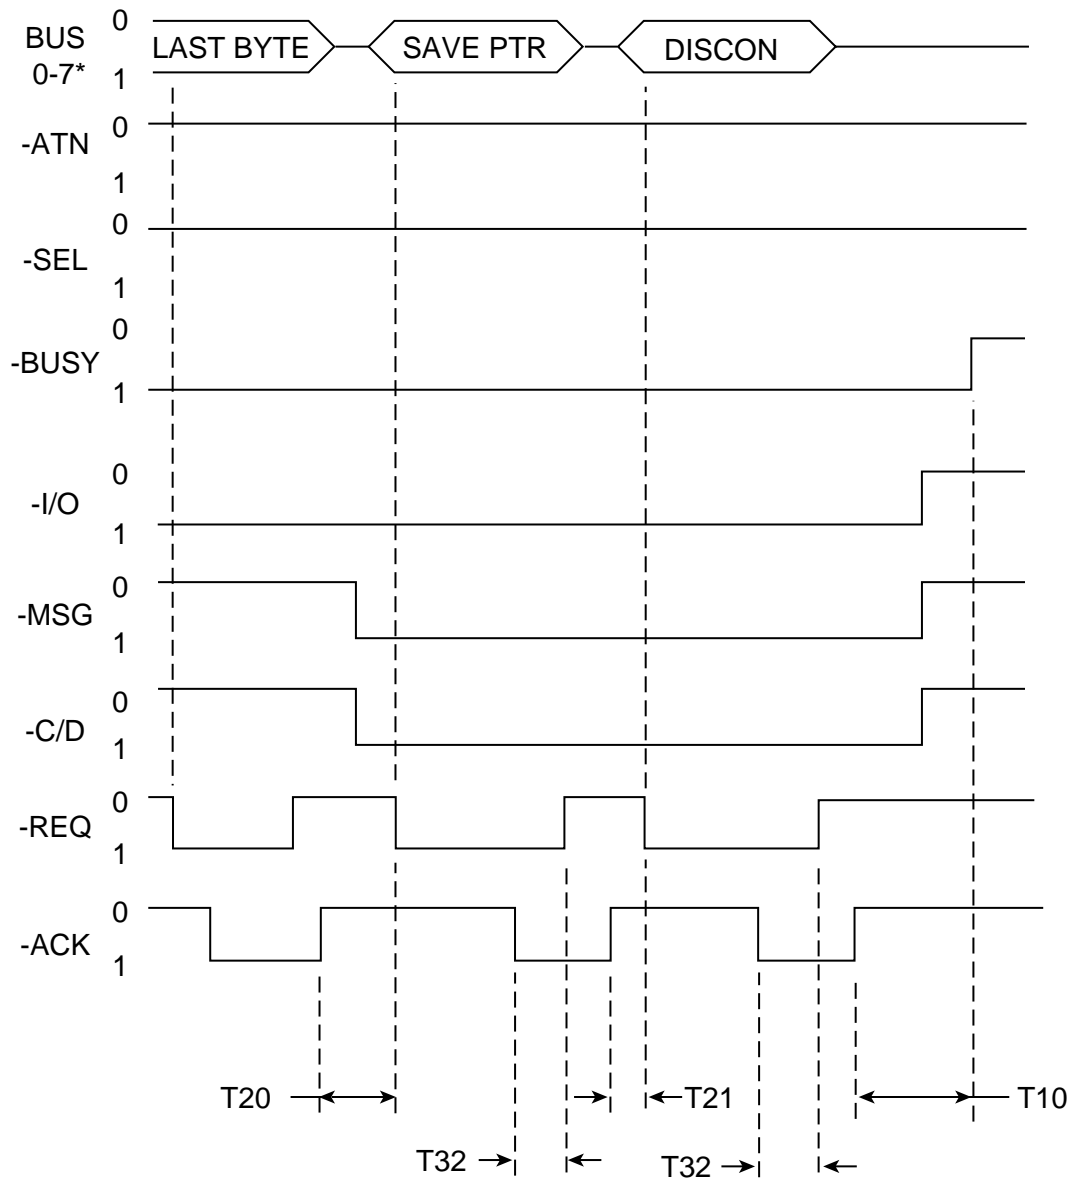


Table 31: Data Out Block Transfer




* and 8-15 if applicable

Table 32: Last Data Byte, Save Pointer Msg, and Disconnect Msg



* and 8-15 if applicable

Table 33: Data in Phase, Status Phase, Command Complete Msg, and Bus Free



* and 8-15 if applicable

4.6 Command processing considerations and exception conditions

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

4.6.1 Auto Contingent Allegiance or Contingent Allegiance

The auto contingent allegiance (NACA=1, see Section 4.2.6) or contingent allegiance (NACA=0) condition shall exist within the task set when the logical unit completes a command by returning a Check Condition status (see Section 4.3).

4.6.1.1 Logical Unit response to Auto Contingent Allegiance or Contingent Allegiance

The auto contingent allegiance (NACA=1, see Section 4.2.6) or contingent allegiance (NACA=0) condition shall not cross task set boundaries and shall be preserved until it is cleared as described in Section 4.6.1.2. If requested by the application client and supported by the protocol and logical unit, sense data shall be returned as described in Section 4.6.4.2.

Notes.

1. The SCSI-2 contingent allegiance condition has had an alternate added and the extended contingent allegiance condition has been replaced in SCSI-3 by auto contingent allegiance in conjunction with the NACA bit.
2. If the SCSI-3 protocol does not enforce synchronization of client/server states, there may be a time delay between the occurrence of the auto contingent allegiance or contingent allegiance condition and the time at which the initiator becomes aware of the condition.

After sending status and a service response of Task Complete, the logical unit shall modify the state of all tasks in the faulted task set as described in ANSI document T10/1157D, clause 7, "Task Set Management."

A task created by the faulted initiator while the auto contingent allegiance condition is in effect may be entered into the faulted task set under the conditions described below.

As described in Section 4.6.1.2, the setting of the NACA bit in the Control byte of the faulting command CDB determines the rules that apply to an ACA or CA condition caused by that command. If the NACA bit was set to zero the SCSI-2 contingent allegiance rules shall apply.

If the NACA bit was set to one in the Control byte of the faulting command, then a new task created by the faulted initiator while the ACA condition is in effect shall not be entered into the faulted task set unless all of the following conditions are true:

- a. The task has the ACA attribute; and
- b. No other task from the faulted initiator having the ACA attribute is in the task set.

If the task is from the faulted initiator and any of the conditions listed above are not met, the newly created task shall not be entered into the task set and shall be completed with a status of ACA Active.

If a task having the ACA attribute is received and no auto contingent allegiance condition is in effect for the task set or if the NACA bit was set to zero in the CDB for the faulting command (i.e., a contingent allegiance condition is in effect), then the ACA task shall be completed with a status of Check Condition. The sense key shall be set to Illegal Request with an additional sense code of Invalid Message Error. As noted in Section 4.6.1.2, any existing contingent allegiance condition shall be cleared and a new auto contingent allegiance (NACA=1) or contingent allegiance (NACA=0) condition shall be established.

The handling of tasks created by initiators other than the faulted initiator depends on the value in the TST field in the Control mode page (see ANSI document SPC-2, T10/1236-D).

If TST=000b, tasks created by other initiators while the ACA or CA condition is in effect shall not be entered into the faulted task set (except for a Persistent Reserve command with a Preempt and Clear action as described in Section 4.6.1.2). Tasks rejected from the task set due to the presence of an ACA or CA condition shall be completed with a status of ACA Active (if NACA=1 in the new command's CDB Control byte, see Section 4.2.6) or Busy (if NACA=0).

If TST=001b, tasks created by one initiator shall not be rejected based on an ACA or CA condition in effect for another initiator. Only ACA or CA condition for the sending initiator (as well as other task set management con-

siderations described in T10/1157D, clause 7, “Task Set Management.”) shall affect acceptance into the task set or rejection for a task from that initiator.

4.6.1.2 Clearing an Auto Contingent Allegiance condition

If the NACA bit is set to zero in the Control byte of the faulting command, then the SCSI-2 rules for clearing contingent allegiance shall apply. In addition, the logical unit shall clear the associated contingent allegiance condition upon sending sense data by means of the autosense mechanism described in Section 4.6.4.2.

While the SCSI-2 rules for clearing the CA condition are in effect, a logical unit that supports the Clear ACA task management function shall ignore all Clear ACA requests and shall return a service response of Function Complete (see ANSI document T10/1157D, Section 6.3).

If the logical unit accepts a value of one for the NACA bit and this bit was set to one in the Control byte of the faulting command, then the SCSI-2 rules for clearing a contingent allegiance condition shall not apply. In this case, the ACA condition shall only be cleared:

- a. As the result of a power on or a logical unit reset (see ANSI document T10/1157D, Section 5.6.1.2);
- b. Through a Clear ACA task management function issued by the faulting initiator as described in T10/1157D;
- c. Through a Preempt and Clear action of a Persistent Reserve Out command that clears the tasks of the faulting initiator (see the SPC-2 standard, T10/1236-D);
- d. A command with the ACA attribute terminates with a Check Condition status. The state of all tasks in the task set when an auto contingent allegiance condition is cleared shall be modified as described in T10/1157D, clause 7, “Task Set Management.”

4.6.2 Overlapped commands

An overlapped command occurs when an application client reuses a Task Address (see Note 3) in a new command before a previous task to which that address was assigned completes its task lifetime as described in T10/1157D, clause 7, “Task and Command Lifetimes.”

Each SCSI-3 protocol standard shall specify whether or not a logical unit is required to detect overlapped commands. A logical unit that detects an overlapped command shall abort all tasks for the initiator in the task set and shall return Check Condition status for that command. If the overlapped command condition was caused by an untagged task or a tagged task with a tag value exceeding FFh, then the sense key shall be set to Aborted Command and the additional sense code shall be set to Overlapped Commands Attempted. Otherwise, an additional sense code of Tagged Overlapped Tasks shall be returned with the additional sense code qualifier byte set to the value of the duplicate tag.

Notes.

1. An overlapped command may be indicative of a serious error and, if not detected, could result in corrupted data. This is considered a catastrophic failure on the part of the initiator. Therefore, vendor-specific error recovery procedures may be required to guarantee the data integrity on the medium. The target logical unit may return additional sense data to aid in this error recovery procedure (e.g., sequential-access devices may return the residue of blocks remaining to be written or read at the time the second command was received).
2. Some logical units may not detect an overlapped command until after the command descriptor block has been received.
3. See Glossary.

4.6.3 Incorrect Logical Unit selection

The target’s response to an incorrect logical unit identifier is described in the following paragraphs.

The logical unit identifier may be incorrect because:

- a. The target does not support the logical unit (e.g., some targets support only one peripheral device). In response to any other command except Request Sense and Inquiry, the target shall terminate the command with Check Condition status. Sense data shall be set to the values specified for the Request Sense command in item b below;

- b. The target supports the logical unit, but the peripheral device is not currently attached to the target. In response to an Inquiry command the target shall return the Inquiry data with the peripheral qualifier set to the value required in the SPC-2 standard. 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.

In response to any other command except Request Sense and Inquiry, the target shall terminate the command with Check Condition status. Sense data shall be set to the values specified for the Request Sense command above;

- c. The target supports the logical unit and the peripheral device is attached, but not operational. In response to an Inquiry command the target shall return the Inquiry data with the peripheral qualifier set to the value required in Section 5.0. In response to Request Sense, the target shall return sense data.

The target's response to any command other than Inquiry and Request Sense is vendor specific; or

- d. The target supports the logical unit but is incapable of determining if the peripheral device is attached or is not operational when it is not ready.

In response to an Inquiry command the target shall return the Inquiry data with the peripheral qualifier set to the value specified in Section 5.0. In response to a Request Sense command the target shall return the Request Sense data with a sense key of No Sense unless an auto contingent allegiance exists. The target's response to any other command is vendor specific.

4.6.4 Sense data

Sense data shall be made available by the logical unit in the event a command completes with a Check Condition status or other conditions. The format, content and conditions under which sense data shall be prepared by the logical unit are specified in this manual, the SPC-2 standard, and applicable SCSI-3 protocol standard.

Sense data shall be preserved by the logical unit for the initiator until it is transferred by one of the methods listed below or until another task from that initiator is entered into the task set.

The sense data may be transferred to the initiator through any of the following methods:

- a. The Request Sense command specified in Section 5.0;
- b. An asynchronous event report; or
- c. Autosense delivery.

The following clauses describe the last two transfer methods.

4.6.4.1 Asynchronous Event Reporting

Asynchronous Event Reporting is used by a logical unit to signal another device that an asynchronous event has occurred. The mechanism automatically returns sense data associated with the event. Support for asynchronous event reporting is a logical unit option.

The control mode page contains parameters affecting the use of asynchronous event reporting (see Section 5.0).

Asynchronous Event Reporting is used to signal a device that one of the four events listed below has occurred:

- a. an error condition was encountered after command completion;
- b. a newly initialized device is available;
- c. some other type of unit attention condition has occurred; or
- d. an asynchronous event has occurred.

An example of the first case above occurs in a device that implements a write cache. If the target is unable to write cache data to the medium, it may use an asynchronous event report to inform the initiator of the failure.

An example of the second case above is a logical unit that generates an asynchronous event report, following a power-on cycle, to notify other SCSI devices that it is ready to accept I/O commands.

Sense data accompanying the report identifies the condition (see Section 4.6.4).

An error condition or unit attention condition shall be reported to a specific initiator once per occurrence of the event causing it. The logical unit may choose to use an asynchronous event report or to return Check Condition

status on a subsequent command, but not both. Notification of an error condition encountered after command completion shall be returned only to the initiator that sent the affected task or tasks.

Asynchronous event reports may be used to notify devices that a system resource has become available. If a logical unit uses this method of reporting, the sense key in the AER sense data shall be set to Unit Attention.

4.6.4.2 Autosense

Autosense is the automatic return of sense data to the application client coincident with the completion of an SCSI-3 command under the conditions described below. The return of sense data in this way is equivalent to an explicit command from the application client requesting sense data immediately after being notified that an ACA condition has occurred. Inclusion of autosense support in an SCSI-3 protocol standard is optional.

As specified in Section 4, the application client may request autosense service for any SCSI command. If supported by the protocol and logical unit and requested by the application client, the device server shall only return sense data in this manner coincident with the completion of a command with a status of Check Condition. After autosense data is sent, the sense data and the CA (NACA=0), if any, shall then be cleared. Autosense shall not affect ACA (NACA=1), see Section 4.6.1.

Protocol standards that support autosense shall require an autosense implementation to:

- a. Notify the logical unit when autosense data has been requested for a command; and
- b. Inform the application client when autosense data has been returned upon command completion (see Section 4.0).

It is not an error for the application client to request the automatic return of sense data when autosense is not supported by the SCSI-3 protocol or logical unit implementation. If the application client requested the return of sense data through the autosense facility and the protocol service layer does not support this feature, then the confirmation returned by the initiator's service delivery port should indicate that no sense data was returned. If the protocol service layer supports autosense but the logical unit does not, then the target should indicate that no sense data was returned. In either case, sense information shall be preserved and the application client may issue a command to retrieve it.

4.6.5 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 non-volatile memory.
8. An event occurs that requires the attention of the initiator.
9. A Clear Queue message is 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.
11. Tasks for this initiator were changed by another initiator.
12. The version or level of microcode has changed.

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 101.

Logical units may queue unit attention conditions. After the first unit attention condition is cleared, another unit attention condition may exist (e.g., a power on condition followed by a microcode change condition).

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 the drive generates the auto contingent allegiance condition), 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 the drive generates the auto contingent allegiance 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 101. 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.

4.6.6 Target Hard Reset

A target hard reset is a target response to a Target Reset task management request (see T10/1157D), or a reset event within the service delivery subsystem. The definition of target reset events is protocol and interconnect specific. Each SCSI-3 product standard shall specify the response to a target reset event including the conditions under which a target hard reset shall be executed.

To execute a hard reset, a target shall initiate a logical unit reset for all attached logical units as described in Section 4.6.7.

4.6.7 Logical Unit Reset

A logical unit reset is a response to a Logical Unit Reset task management request (see T10/1157D), or some other logical unit reset event, such as a target hard reset (see Section 4.6.6). The definition of such events may be device-specific or dependent on the protocol and interconnect. Each appropriate SCSI-3 standard shall specify the conditions under which a logical unit reset shall be executed.

To execute a logical unit reset the logical unit shall:

- a. Abort all tasks in its task set;
- b. Clear an auto contingent allegiance condition, if one is present;
- c. Release all SCSI device reservations;
- d. Return the device's operating mode to the appropriate initial conditions, similar to those conditions that would be found following device power-on. The Mode Select conditions (see the SPC standard) shall be restored to their saved values if saved values have been established. Mode Select conditions for which no saved values have been established shall be returned to their default values;
- e. set a Unit Attention condition (see Section 4.6.5).
- f. Initiate a logical unit reset for all nested logical units.

In addition to the above, the logical unit shall execute any additional functions required by the applicable standards.

4.7 Queued tasks (formerly “Queued I/O Processes”)

Queuing of tasks 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 task queuing allows the drive to accept multiple commands from each initiator. Untagged task 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 task queuing mode if the initiator does not send task queue tag messages.

Initiators may add or delete commands to the queue for the 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 task queuing

Untagged task 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 task may be initiated any time the Bus Free phase exists even if another task from a different initiator is being executed. If the disconnect privilege is not granted, the drive returns Busy status to the new task.

The I T L nexus specifies the relationship so that the drive can always reconnect to the initiator to restore the pointers for task 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 task queuing

Tagged task queuing allows a drive to accept multiple commands from the same or different initiators until the drive's task queue is full. A new task 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 task.

The Task Queue Tag messages (see Table 5) allow the initiator to establish a unique I T L Q nexus to identify each task. Each task 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 task and the initiator to restore the correct set of pointers. An initiator may have several task ongoing to the same or different logical unit as long as each has a unique nexus.

If only Simple Task 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 96). 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 Task Queue Tag when reconnecting to the initiator.

If Ordered Task Queue Tag messages are used, the drive executes the commands in the order received with respect to other commands received with Ordered Task Queue Tag messages. All commands received with a Simple Task Queue Tag message prior to a command received with an Ordered Task Queue Tag message, regardless of initiator, are executed before that command with the Ordered Task Queue Tag message. All commands received with a Simple Task Queue Tag message after a command received with an Ordered Task Queue Tag message, regardless of initiator, are executed after that command with the Ordered Task Queue Tag message.

A command received with a Head of Task Queue Tag message is placed first in the queue, to be executed next. A command received with a Head of Task Queue Tag message does not suspend a task for which the drive has begun execution. Consecutive commands received with Head of Task 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 Table 96).

A task received from an initiator without a task 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. A task received from an initiator with a task 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 Task Queue Tag message. Use of the Head of Task 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 Task 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 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 target 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 tasks that are already in progress, the target may disconnect to allow completion of any tasks 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.

4.10 Incorrect initiator connection

An incorrect initiator connection occurs on a reconnection if:

1. an initiator attempts to reconnect to a task, and
2. a soft reset condition has not occurred, and
3. the initiator does not send an Abort, Abort Tag, Bus Device Reset, Clear Task Set, Continue Task, or Terminate Task 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 tasks 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 a task 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.

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.

Table 34 lists commands alphabetically by command name. Table 35 lists commands by operation code.

Table 34. Commands sorted by command name

Command name	Command operation code	Command type		Length			Reference
		All devices	Direct access devices	6 bytes	10 bytes	12 bytes	
Change Definition	40h	x					5.1 (see page 78)
Compare	39h	x					5.2 (see page 79)
Copy	18h						5.3 (see page 79)
Copy and Verify	3Ah	x					5.4 (see page 79)
Format Unit	04h		x	x			5.5 (see page 80)
Inquiry	12h	x		x			5.6 (see page 87)
Lock-unlock Cache	36h						5.7 (see page 101)
Log Select	4Ch	x			x		5.8 (see page 102)
Log Sense	4Dh	x			x		5.9 (see page 110)
Mode Select (6)	15h		x	x			5.10 (see page 112)
Mode Select (10)	55h		x		x		5.11 (see page 116)
Mode Sense (6)	1Ah		x	x			5.12 (see page 119)
Mode Sense (10)	5Ah		x		x		5.13 (see page 149)
Move Medium	A7h						5.14 (see page 152)
Persistent Reserve In	5Eh		x				5.15 (see page 153)
Persistent Reserve Out	5Fh		x				5.16 (see page 158)
Prefetch	34h		x				5.17 (see page 161)
Prevent/Allow Medium Removal	1Eh						5.18 (see page 161)
Read	08h		x	x			5.19 (see page 162)
Read Buffer	3Ch	x			x		5.20 (see page 164)
Read Capacity	25h		x		x		5.21 (see page 167)
Read Defect Data (10)	37h		x		x		5.22 (see page 169)
Read Defect Data (12)	B7h		x			x	5.23 (see page 171)
Read Element Status	B4h						5.24 (see page 172)
Read Extended	28h		x		x		5.25 (see page 173)
Read Long	3Eh		x		x		5.26 (see page 175)
Reassign Blocks	07h		x	x			5.27 (see page 176)
Rebuild	81h		x		x		5.28 (see page 178)
Receive Diagnostic Results	1Ch	x		x			5.29 (see page 181)
Regenerate	82h		x				5.30 (see page 185)
Release (6)	17h		x	x			5.31 (see page 187)
Release (10)	57h		x		x		5.32 (see page 188)
Report LUNs	A0h		x				5.33 (see page 189)
Request Sense	03h	x		x			5.34 (see page 191)
Reserve (6)	16h		x	x			5.35 (see page 200)
Reserve (10)	56h		x		x		5.36 (see page 202)
Rezero Unit	01h		x	x			5.37 (see page 203)
Search Data Equal	31h						5.38 (see page 203)
Search Data High	30h						5.39 (see page 203)
Search Data Low	32h						5.40 (see page 203)
Seek	0Bh		x	x			5.41 (see page 204)
Seek Extended	2Bh		x		x		5.42 (see page 205)
Send Diagnostic	1Dh	x		x			5.43 (see page 206)
Set Limits	33h						5.44 (see page 208)
Start/Stop Unit	1Bh		x	x			5.45 (see page 209)
Synchronize Cache	35h		x		x		5.46 (see page 210)
Test Unit Ready	00h	x		x			5.47 (see page 211)
Verify	2Fh		x		x		5.48 (see page 212)
Write	0Ah		x	x			5.49 (see page 213)
Write and Verify	2Eh		x		x		5.50 (see page 214)
Write Buffer	3Bh	x			x		5.51 (see page 215)
Write Extended	2Ah		x		x		5.52 (see page 219)
Write Long	3Fh		x		x		5.53 (see page 221)
Write Same	41h		x				5.54 (see page 222)
XDRead	52h		x				5.55 (see page 223)
XDWrite	50h		x				5.56 (see page 225)
XDWrite Extended	80h		x				5.57 (see page 226)
XPWrite	51h		x				5.58 (see page 228)

Table 35. Commands sorted by operation code

Command operation code	Command name	Command type		Length			Reference
		All devices	Direct access devices	6 bytes	10 bytes	12 bytes	
00h	Test Unit Ready	x		x			5.47 (see page 211)
01h	Rezero Unit		x	x			5.37 (see page 203)
03h	Request Sense	x		x			5.34 (see page 191)
04h	Format Unit		x	x			5.5 (see page 80)
07h	Reassign Blocks		x	x			5.27 (see page 176)
08h	Read		x	x			5.19 (see page 162)
0Ah	Write		x	x			5.49 (see page 213)
0Bh	Seek		x	x			5.41 (see page 204)
12h	Inquiry	x		x			5.6 (see page 87)
15h	Mode Select (6)		x	x			5.10 (see page 112)
16h	Reserve (6)		x	x			5.35 (see page 200)
17h	Release (6)		x	x			5.31 (see page 187)
18h	Copy						5.3 (see page 79)
1Ah	Mode Sense (6)		x	x			5.12 (see page 119)
1Bh	Start/Stop Unit		x	x			5.45 (see page 209)
1Ch	Receive Diagnostic Results	x		x			5.29 (see page 181)
1Dh	Send Diagnostic	x		x			5.43 (see page 206)
1Eh	Prevent/Allow Medium Removal						5.18 (see page 161)
25h	Read Capacity		x		x		5.21 (see page 167)
28h	Read Extended		x		x		5.25 (see page 173)
2Ah	Write Extended		x		x		5.52 (see page 219)
2Bh	Seek Extended		x		x		5.42 (see page 205)
2Eh	Write and Verify		x		x		5.50 (see page 214)
2Fh	Verify		x		x		5.48 (see page 212)
30h	Search Data High						5.39 (see page 203)
31h	Search Data Equal						5.38 (see page 203)
32h	Search Data Low						5.40 (see page 203)
33h	Set Limits						5.44 (see page 208)
34h	Prefetch		x				5.17 (see page 161)
35h	Synchronize Cache		x		x		5.46 (see page 210)
36h	Lock-unlock Cache						5.7 (see page 101)
37h	Read Defect Data (10)		x		x		5.22 (see page 169)
39h	Compare	x					5.2 (see page 79)
3Ah	Copy and Verify	x					5.4 (see page 79)
3Bh	Write Buffer	x			x		5.51 (see page 215)
3Ch	Read Buffer	x			x		5.20 (see page 164)
3Eh	Read Long		x		x		5.26 (see page 175)
3Fh	Write Long		x		x		5.53 (see page 221)
40h	Change Definition	x					5.1 (see page 78)
41h	Write Same		x				5.54 (see page 222)
4Ch	Log Select	x			x		5.8 (see page 102)
4Dh	Log Sense	x			x		5.9 (see page 110)
50h	XDWrite		x				5.56 (see page 225)
51h	XPWrite		x				5.58 (see page 228)
52h	XDRead		x				5.55 (see page 223)
55h	Mode Select (10)		x		x		5.11 (see page 116)
56h	Reserve (10)		x		x		5.36 (see page 202)
57h	Release (10)		x		x		5.32 (see page 188)
5Ah	Mode Sense (10)		x		x		5.13 (see page 149)
5Eh	Persistent Reserve In		x				5.15 (see page 153)
5Fh	Persistent Reserve Out		x				5.16 (see page 158)
80h	XDWrite Extended		x				5.57 (see page 226)
81h	Rebuild		x		x		5.28 (see page 178)
82h	Regenerate		x				5.30 (see page 185)
A0h	Report LUNs		x				5.33 (see page 189)
A7h	Move Medium						5.14 (see page 152)
B7h	Read Defect Data (12)		x			x	5.23 (see page 171)
B4h	Read Element Status						5.24 (see page 172)

5.1 Change Definition command (40h)

The Change Definition command (Table 36) 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.

Table 36: Change Definition command (40h)

Bit Byte	7	6	5	4	3	2	1	0
0	0	1	0	0	0	0	0	0
1	Logical Unit Number [1]			Reserved				
2	Reserved						Save [2]	
3	Reserved	Definition Parameter [3]						
4	Reserved							
5	Reserved							
6	Reserved							
7	Reserved							
8	Parameter Data Length [4]							
9	Control [5]							

Notes.

- [1] The Logical Unit Number must be zero.
- [2] A Save Control (Save) bit 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 37.

Table 37: Definition Parameter field

Value	Meaning of Definition Parameter
00h	Use Current Operation 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] Use of the Control byte by the drive is a customer unique option. See paragraph 4.2.6.

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 55).

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.

5.2 Compare command (39h)

Not implemented. If received, a Check Condition status is sent.

5.3 Copy command (18h)

Not implemented. If received, the drive sends a Check Condition status and a Sense Key of Illegal Request.

5.4 Copy and Verify command (3Ah)

Not implemented. If received, a Check Condition status is sent.

5.5 Format Unit command (04h)

The Format Unit command (Table 38) 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 39.

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.

Table 38: Format Unit command (04h)

Bit Byte	7	6	5	4	3	2	1	0
0	0	0	0	0	0	1	0	0
1	Logical Unit Number [1]			FmtData [2]	CmpLst [2]	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]

Notes.

- [1] The Logical Unit Number must be zero.
- [2] A Format Data (FmtData) bit of one indicates that the Format Unit Parameter list (Table 40) is supplied during the Data Out phase. The Data Out phase consists of a defect list header (Table 41), followed by an initialization pattern descriptor (Table 42) (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 41).
- The Defect List Format field specifies additional information related to the defect list (see Table 39 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 5.19). Interleave values other than zero or one are vendor specific.
- [4] See “Control byte,” paragraph 4.2.6.

5.5.1 Format Unit parameter definition

The following definitions of flaw categories are supplied to help in understanding the alternatives listed in Table 39.

- P Primary Defect (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 (C type) flawed sectors are sectors that fail a format verify during the format function.
- D Data Defect (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 (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 39: Format Unit parameter definition (format variations)

CDB-Byte (see Table 38)					Comments
FMT DATA	CMP LIST	Defect List Format			
Bit 4	Bit 3	Bit 2	Bit 1	Bit 0	
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 [2]					
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 drive's current G list [1].
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 [1].
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 46). The drive erases any previous G list [1].
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 46). The drive also reallocates all sectors in the drive's current G list [1].
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 47). The drive erases any previous G list [1].
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 47). The drive also reallocates all sectors in the drive's current G list [1].

Notes.

- [1] Byte one of the Defect List Header determines whether the P and C defects are reallocated (see Table 41).
- [2] See individual drive's Product Manual, Volume 1, for support/nonsupport.

5.5.2 Format Unit parameter list

The defect list shown in Table 41 contains a four byte header followed by one or more defect descriptors (Table 45). The Defect List Length in each table specifies the total length in bytes of the defect descriptors that follow. In Table 41 the Defect List Length is equal to eight times the number of defect descriptors.

Table 40: Format Unit parameter list

Bit Byte	7	6	5	4	3	2	1	0
[1]	Defect List Header							
[2]	Initialization Pattern Descriptor							
[3]	Defect Descriptor(s)							

[1] See Table 41.

[2] See Table 42.

[3] See Tables 45, 46, and 47.

5.5.2.1 Defect List header

Table 41: Defect List header

Bit Byte	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	Defect List Length [9]							
3								

Notes.

[1] Format Options Valid (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] Disable Primary (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's 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] Disable Certification (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 0 when the user sets the IP bit (see Table 41) and specifies a format pattern of greater than 1 byte (see Table 42), and/or if the IP Modifier bits are set to 01 or 10 (see Table 42). 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] Stop Format (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.

- [5] Initialization Pattern (IP). An 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 44). 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 43.
- [6] Disable Saving Parameters (DSP). A 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 (Immed). An 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's Product Manual, Volume 1, for applicability of this feature.
- [8] Vendor Specific (VS). 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 46 or Table 47.

Note. The initiator may not use any previously defined "C," "G," or "D" lists if the sector size (block length) has been changed.

5.5.2.2 Initialization Pattern Descriptor

The Initialization Pattern option specifies that the logical blocks contain the specified initialization pattern. The Initialization Pattern Descriptor (Table 42) is sent to the target as part of the Format Unit parameter list.

Table 42: Initialization Pattern Descriptor

Bit Byte	7	6	5	4	3	2	1	0
0	IP Modifier [1]		Reserved					
1	Pattern Type [2]							
2	(MSB)							
3	Initialization Pattern Length [3]						(LSB)	
0-n	Initialization Pattern [1]							

Notes.

- [1] Initialization Pattern (IP) and IP Modifier. The Initialization Pattern is modified in accordance with the specification of the IP Modifier field given in Table 43.

Table 43: Initialization Pattern and IP Modifier

IP Modifier	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.

[2] Pattern Type. Table 44 gives pattern type information.

Table 44: Initialization Pattern Type

Pattern Type	Description
00h	Use default pattern. If the initialization pattern length is not zero, this is an error and 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.
01h	Use pattern supplied by host computer. If the initialization pattern length is zero, this is an error and 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.
02h–7Fh	Reserved.
80h–FFh	Vendor specific.

[3] The Initialization Pattern Length field indicates the number of bytes contained in the initialization pattern. If the length exceeds the current logical block size this is an error and 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. The initialization pattern is sent by the host for the drive to write in each logical block by the format command.

5.5.2.3 Defect List formats

This section describes the format of the defect list that follows the Defect List Header described in Table 41. Three formats are possible:

- Block format (Seagate drives do not support the block format, except as a customer special)
- Bytes from Index format (see Table 46)
- Physical Sector format (see Table 47)

Table 45: Defect Descriptors

0-n	Defect Descriptor Bytes (Refer to Tables 46 and 47 for formats)
-----	--

Table 46: 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 Defect
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 38 and 39).

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 47: Defect Descriptor Bytes—Physical Sector 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 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 38 and 39.

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.

5.6 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 5.6.2. Several Inquiry commands may be sent to request the vital product data pages instead of the standard data shown in Table 49.

Table 48: Inquiry command (12h)

Bit Byte	7	6	5	4	3	2	1	0
0	0	0	0	1	0	0	1	0
1	Logical Unit Number [1]			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 [5]	Link [5]

Notes.

- [1] The Logical Unit Number 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 non-volatile 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.

5.6.1 Drive Standard 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's Product Manual, Volume 1). The standard Inquiry data is given in Table 49.

Table 49: Drive Standard Inquiry Data Format

Bit Byte	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 Version [2]		
3	0 AERC [3]	0 TrmTsk [4]	0 NormACA [5]	0 HiSupport [6]	Response Data Format [7]			
4	Additional Length [8]							
5	SCCS [9]	Reserved						
6	BQue [10]	EncServ [11]	Vendor Specific [12]	MultiP [13]	MChangr [14]	AckReqQ [15]	ADR32 [16]	ADR16 [17]
7	RelAdr [18]	WBus32 [19]	WBus16 [20]	Sync [21]	Linked [22]	TrnDis [23]	CmdQue [24]	Soft Reset [25]
8 : 15	Vendor Identification [26]							
16 : 31	Product Identification [27]							
32 : 35	Product Revision Level [28]							
36 : 43	Drive Serial Number [29]							
44 : 55	Unused Vendor Specific Area (00h)							
56 : 95	Reserved (00h) [30]							
96 : 143	Copyright Notice [31]							
144 : 147	Distribution Diskette Serial Number [32]							

Notes.

- [1] The Peripheral Qualifier field 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 50.

Table 50: ANSI-approved version

Code	Description
0h	The device does not claim conformance to any standard.
1h	Obsolete (formerly SCSI-1).
2h	The device complies to ANSI X3.131-1994 (SCSI-2).
3h	The device complies to ANSI X3.301-1997.
4h	The device complies to ANSI T10/1236.
5h-7h	Reserved.
[3]	An Asynchronous Event Reporting Capability (AERC) bit of zero indicates that the drive does not support the asynchronous event reporting capability.
[4]	A Terminate Task (TrmTsk) bit of zero indicates that the drive does not support the Terminate Task management function.
[5]	The Normal ACA Supported (NormACA) bit of one indicates that the device server supports setting the NACA bit to one in the Control Byte of the CDB (as defined in ANSI document SCSI-3 Architectural Model-2, T10/1157). A NormACA bit of zero indicates that the device server does not support setting the NACA bit to one (see paragraph 4.2.6, "Control field.")
[6]	A Hierarchical Support (HiSupport) bit of zero indicates the target does not use the hierarchical addressing model to assign LUNs to logical units. A HiSupport bit of one indicates the target uses the hierarchical addressing model to assign LUNs to logical units. When the HiSupport bit is one, the device server shall support the Report LUNs command (see Section 5.33).
[7]	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 Controller Command Set-2 standard. 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.
[8]	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.
[9]	An SCC Supported (SCCS) bit of one indicates that the device contains an embedded storage array controller component. See SCC-2 (SCSI-3 Controller Command Set-2, T10/1225-D) for details about storage array controller devices. An SCCS bit of zero indicates that the device does not contain an embedded storage array controller component.
[10]	The Basic Queuing (BQue) bit shall be zero if the CmdQue (see note [24]) bit is one. When the CmdQue bit is zero, the BQue bit shall have the following meaning. A BQue bit of zero indicates that the device does not support tagged tasks (command queuing) for this logical unit. A value of one indicates that the device supports, for this logical unit, the basic task management model defined by ANSI document SCSI-3 Architectural Model-2, T10/1157.
[11]	An Enclosure Services (EncServ) bit of one indicates that the device contains an embedded enclosure services component. See ANSI Document SCSI-3 Enclosure Services Command Set, NCITS 305-199X, T10/1212-D for details about enclosure services, including a device model for an embedded enclosure services device. An EncServ bit of zero indicates that the device does not contain an embedded enclosure services component.
[12]	Vendor Specific. Dual port A-B on some drives.
[13]	A Multi Port (MultiP) bit of one shall indicate that this is a multi-port (2 or more ports) device and conforms to the SCSI-3 multi-port device requirements found in the applicable standards. A value of zero indicates that this device has a single port and does not implement the multi-port requirements.
[14]	A Medium Changer (MChngr) bit of one indicates that the device is embedded within or attached to a medium transport element. See SMC (T10/999-D) for details about medium changers, including a device model for an attached medium changer device. The MChngr bit is valid only when the RMB bit is equal to one. An MChngr bit of zero indicates that the device is not embedded within or attached to a medium transport element.
[15]	An ACKQ/REQQ (AckReqQ) bit of one indicates that the target supports a request and acknowledge data transfer handshake on the secondary bus.

- [16] An ADR32 bit of one indicates the drive supports 32-bit wide SCSI addresses. Zero indicates non-support.
 - [17] An ADR16 bit of one indicates the drive supports 16-bit wide SCSI addresses. Zero indicates non-support. A zero in both bits [16] and [17] indicates the drive supports only 8-bit SCSI addresses.
 - [18] A Relative Address (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.
 - [19] 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.
 - [20] 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 WBus16 and WBus32 bits are zero, the drive only supports 8-bit wide data transfers.
- [21] 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.
 - [22] 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.
 - [23] When the Transfer Disable (TrnDis) bit is one, it indicates that the drive supports the Continue Task and Target Transfer Disable messages. When this bit is zero, it indicates the drive does not support the above two messages.
 - [24] A Command Queuing (CmdQue) bit of one indicates that the device supports tagged tasks (command queuing) for this logical unit (see ANSI document SCSI-3 Architectural Model-2, T10/1157). A value of zero indicates the device server may support tagged tasks for this logical unit (see note [10]). Table 51 summarizes the relationship of the BQue and CmdQue bits.

Table 51: Relationship of BQue and CmdQue bits

BQue	CmdQue	Description
0	0	No command queuing of any kind supported.
0	1	Command queuing with all types of task tags supported.
1	0	Basic command queuing (limited types of task tags) supported.
1	1	Illegal combination of BQue and CmdQue bits.

- [25] A 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).
- [26] The Vendor Identification field contains the ASCII data giving vendor name ("SEAGATE").
- [27] The Product Identification field contains ASCII data giving the drive model number. The data shall be left-aligned within this field.
- [28] 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.
- [29] The Drive Serial Number field contains the 8 bytes of ASCII data "XXXXXXXX," where the value XXXXXXXX is the drive serial number.
- [30] The reserved area from byte 56 through byte 95 is filled with 00h.
- [31] The Copyright Notice field contains the 48 bytes of ASCII data "Copyright (c) 199X Seagate All rights reserved," where "X" indicates the current year.
- [32] The Distribution Diskette Serial Number field contains the 4bytes of ASCII data "XXXX," where the value XXXX is the Distribution Diskette serial number.

5.6.2 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 52: Supported Vital Product Data pages (00h)

Bit Byte	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 Peripheral Device Type field value of 00h indicates a direct-access device (magnetic disc) is connected to this logical unit.
- [2] Page 00h 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.
- [4] The Supported Page List field contains a list of all vital product data page codes implemented for the drive in ascending order beginning with page code 00h. See Table 53.

Table 53: Vital Product Data Page Codes

Page code	Description
00h	Supported Vital Product Data pages
80h	Unit Serial Number page
81h	Implemented Operating Definition page
83h	Device Identification 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)
D1h	Vendor Specific page
D2h	Vendor Specific page
D3h	Vendor Specific page
D4h	Vendor Specific page

5.6.3 Unit Serial Number page (80h)

Table 54: Unit Serial Number page (80h)

Bit Byte	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 [3]							
4 : 11	Product Serial Number [4]							
12 : 23	Board Serial Number [5]							

Notes.

- [1] The Peripheral Qualifier field and Peripheral Device Type field value of 00h indicates a direct-access device (magnetic disc) is connected to this logical unit.
- [2] Page 80h provides the product serial number and product circuit board number for the drive.
- [3] The Page Length field specifies the length in bytes of the product serial number page. Older products that only support the Product Serial Number parameter will have a page length of 08h, while newer products that support both parameters will have a page length of 14h. 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.

5.6.4 Implemented Operating Definition page (81h)

Table 55: Implemented Operating Definition page (81h)

Bit Byte	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 (n-3) [3]							
4	0	Current Operating Definition [4]						
5 6	SAVIMP [5]	Default Operating Definition [5]						
n	SAVIMP [5]	Supported Operating Definition List [5]						

Notes.

- [1] The Peripheral Qualifier field and Peripheral Device Type field value of 00h indicates a direct-access device (magnetic disc) is connected to this logical unit.

- [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 36).
- [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 56).

Table 56: 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

- [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.

5.6.5 Device Identification page (83h)

The Device Identification page (Table 57) provides the means to retrieve zero or more identification descriptors applying to the logical unit. Logical units may have more than one identification descriptor (e.g., if several types or associations of identifier are supported).

Device identifiers, if any, shall be assigned to the peripheral device (e.g., a disk drive) and not to the currently mounted media, in the case of removable media devices. Media identification is outside the scope of this manual. Operating systems are expected to use the device identifiers during system configuration activities to determine whether alternate paths exist for the same peripheral device.

Note. In the case of virtual logical units (e.g., volume sets as defined by SCC, SCSI-3 Controller Commands), the Identifier field (see Table 58) should be an IEEE Registered Extended name. The IEEE Registered Extended name has a code of 3h in the Identifier type field and an NAA value of 0110b as defined in FC-PH.

Table 57: Device Identification page (83h)

Bit Byte	7	6	5	4	3	2	1	0
0	Peripheral Qualifier [1]			Peripheral Device Type [1]				
1	Page Code (83h)							
2	Reserved							
3	Page Length (n-3)							
	Identification Descriptor List [2]							
4	Identification Descriptor (First)							
	.							
	.							
	.							
n	Identification Descriptor (Last)							

Notes.

- [1] The Peripheral Qualifier field and Peripheral Device Type field value of 00h indicates a direct-access device (magnetic disc) is connected to this logical unit.
- [2] The Identification Descriptor List is formatted per Table 58. Each Identification descriptor contains information identifying the logical unit. If the logical unit is accessible through any other path, it shall return the same identification.

Table 58: Identification Descriptor

Bit Byte	7	6	5	4	3	2	1	0
0	Reserved			Code Set [1]				
1	Reserved		Association [2]		Identifier Type [3]			
2	Reserved							
3	Identifier Length (n-3) [4]							
4	(MSB) Identifier [5]							
n	(LSB)							

Notes.

- [1] The Code Set field specifies the code set used for the identifier field, as described in Table 59. This field is intended to be an aid to software that displays the identifier field.

Table 59: Code Set

Value	Description
0h	Reserved.
1h	The identifier field shall contain binary values.
2h	The identifier field shall contain ASCII graphic codes (i.e., code values 20h through 7Eh).
3h-Fh	Reserved.

- [2] The Association field specifies the entity with which the Identifier field is associated, as described in Table 60.

Table 60: Association

Value	Description
0h	The Identifier field is associated with the addressed physical or logical device.
1h	The identifier field is associated with the port that received the request.
2h-Fh	Reserved.

- [3] The Identifier Type field specifies the format and assignment authority for the identifier, as described in Table 61.

Table 61: Identifier Type

Value	Description
0h	No assignment authority was used and consequently there is no guarantee that the identifier is globally unique (i.e., the identifier is vendor specific).
1h	The first 8 bytes of the identifier field are a Vendor ID (see T10/1236-D, annex C). The organization associated with the Vendor ID is responsible for ensuring that the remainder of the identifier field is unique. One recommended method of constructing the remainder of the identifier field is to concatenate the product identification field from the standard Inquiry data field and the product serial number field from the unit serial number page.
2h	The identifier field contains a Canonical form IEEE Extended Unique Identifier, 64-bit (EUI-64). In this case, the identifier length field shall be set to 8. Note that the IEEE guidelines for EUI-64 specify a method for unambiguously encapsulating an IEEE 48-bit identifier within an EUI-64.
3h	The identifier field contains an FC-PH Name_Identifier. Any FC-PH identifier may be used, including one of the four based on a Canonical form IEEE company_id.identifier field
4h-Fh	Reserved.

- [4] The Identifier Length field specifies the length in bytes of the Identifier. If the Allocation Length field of the command descriptor block is too small to transfer all of the identifier, the Identifier length shall not be adjusted to reflect the truncation.
- [5] The Identifier field contains the identifier as described by the Identifier Type, Code Set, and Identifier Length fields.

The example described in this paragraph and shown in Table 62 is not a normative part of this standard. This example of a complete device identification VPD page assumes that the product is a direct-access device with a T10 Vendor ID of "XYZ_Corp," a product identification of "Super Turbo Disk," and a product serial number of "2034589345." Furthermore, it is assumed that the manufacturer has been assigned a 24-bit IEEE company_id of 01ABCDh by the IEEE Registration Authority Committee and that the manufacturer has assigned a 24-bit extension_identifier of 234567h to this logical unit. The combined 48-bit identifier is reported in the 64-bit format as defined by the IEEE 64-bit Global Identifier (EUI-64) standard. The data returned in the device identification VPD page for this logical unit is shown in Table 62.

Table 62: Device identification page example

Bytes	Hexadecimal values	ASCII values
00–15	00 83 00 32 02 01 00 22 58 59 5A 5F 43 6F 72 70	...2...XYZ_Corp
16–31	53 75 70 65 72 20 54 75 72 62 6F 20 44 69 73 6B	Super Turbo Disk
32–47	32 30 33 34 35 38 39 33 34 35 01 02 00 08 01 AB	2034589345.....
48–53	CD FF FF 23 45 67
Notes:	a) Non-printing ASCII characters are shown as periods (.). b) Byte 00 is the beginning of the VPD page (see Table 57). c) Byte 04 is the beginning of the Identification descriptor for the Vendor ID based identifier (Identifier type 1, see Table 61). d) Byte 42 is the beginning of the Identification descriptor for the EUI-64 identifier (Identifier type 2, see Table 61).	

5.6.6 Firmware Numbers page (C0h)

Table 63 applies to drives that do not supply a servo RAM release number and date and a ROM release date. Table 64 applies to model families other than those covered by Table 63.

Table 63: Firmware Numbers page (C0h)

Bit Byte	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 (18h) [3]							
4 : 11	SCSI Firmware Release Number [4]							
12 : 19	ASCII space characters							
20 : 27	Servo ROM Release Number [4] [5]							

Notes.

- [1] The Peripheral Qualifier field and Peripheral Device Type field value of 00h indicates a direct-access device (magnetic disc) is connected to this logical unit.
- [2] Page C0h provides the firmware release numbers for the drive (see Table 64).
- [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.
- [5] For drives that return only 28 bytes.

Table 64: Firmware Numbers page (C0h) (Applies to model families [5][9] other than those covered by Table 63)

Bit Byte	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 [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 [5]							
32 : 35	Servo ROM Release Date In ASCII [5]							
36 : 43	SAP Firmware Release Number [6] [9]							
44 : 47	SAP Firmware Release Date [6] [9]							
48 : 51	SAP Firmware Release Year [6] [9]							
52 : 55	SAP Manufacturing Key [7] [9]							
56 : 59	Servo Firmware Product Family and Product Family Member IDs [8] [9]							

Notes.

- [1] The Peripheral Qualifier field and Peripheral Device Type field value of 00h indicates a direct-access device (magnetic disc) is connected to this logical unit.
- [2] Page C0h provides the firmware release numbers for the drive.
- [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.
- [5] Applies to model families that return 36 bytes rather than 28. Servo firmware RAM and ROM release dates are added in some models. See "Vital Product Data pages supported" table in the Product Manual, Volume 1, that applies to the drive in question.
- [6] Servo Adaptive Parameters (SAP) are created independently from the servo firmware. Thus, they have their own release information. The SAP Firmware Release Number field contains ASCII data.

- [7] The SAP Manufacturing Key field contains a binary code used by manufacturing to identify the SAP.
- [8] The high order word of the Servo Firmware Product Family and Product Family Member IDs field contains a binary code which corresponds to the Servo Firmware Family ID. The low order word of this field contains a binary code which corresponds to the Product Family Member ID.
- [9] Applies to model families that return 59 bytes rather than 35 or 28. Servo firmware RAM and ROM release dates are added, plus SAP firmware. See “Vital Product Data pages supported” table in the Product Manual, Volume 1, that applies to the drive in question.

5.6.7 Date Code page (C1h)

Table 65: ETF Log Date Code page (C1h)

Bit Byte	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 (Ch) [3]							
4	(MSB)							
:	ETF Log Date in ASCII [4]							
11								
12	(MSB)							
:	Compile Date Code [5]							
19								

Notes.

- [1] The Peripheral Qualifier field and Peripheral Device Type field value of 00h indicates a direct-access device (magnetic disc) is connected to this logical unit.
- [2] Page C1h provides the date code from the drive defect list.
- [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 MMDDYYYY. Leading ASCII zero characters are added to single-digit months or days. [6]
- [5] The Compile Date Code field contains 6 ASCII bytes of data for a date of the form MMDDYYYY. [6]
- [6] Applies to Barracuda 18XL, Cheetah 18XL, and later products.

5.6.8 Jumper Settings page (C2h)

Table 66: Jumper Settings page (C2h)

Bit Byte	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	Reserved							Term Enable [9] [10]

Notes.

- [1] The Peripheral Qualifier field and Peripheral Device Type field value of 00h indicates a direct-access device (magnetic disc) is connected to this logical unit.
- [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] The Delayed Motor Start (DS) bit when set to 1 indicates that this jumper is on.
- [5] The Motor Start (MS) bit when set to 1 indicates that the jumper is on.
- [6] The Write Protect (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 individual drive's Product Manual, Volume 1, or Installation Guide for information on how the drive of interest is set up.

[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 individual drive's Product Manual, Volume 1, for applicability.

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

5.6.9 Device Behavior page (C3h)

The Device Behavior page (Table 67) will be used by the regression tests to determine what behavior should be expected from a particular firmware package.

Table 67: Device Behavior page (C3h)

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

Notes.

- [1] The Peripheral Qualifier field and Peripheral Device Type field value of 00h indicates a direct-access device (magnetic disc) is connected to this logical unit.
- [2] 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.
- [3] 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.
- [4] The Behavior Code and Behavior Code Version Number are jointly assigned by the Firmware Engineering Managers of all SCSI design locations.
- [5] The ASCII Family Number field contains the drive model number. The data is left-aligned within this field.
- [6] 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 38). 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 92).
- [7] The Default Number of Cache Segments field gives the number of segments into which the host requests the drive divide the cache.

5.7 Lock-unlock Cache command (36h)

Not implemented. If received, a Check Condition status is sent.

5.8 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 68. 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 70. These Log Pages contain parameters that command the drive to change selected threshold, or cumulative values of any or all drive logs.

The following tables in Section 5.8 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 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. Section 5.9 describes the Log Sense command, but the tables of this section that apply are not repeated there.

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 would 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 4.6.5, Unit attention condition.

Table 68: Log Select command (4Ch)

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

[1] The Logical Unit Number must be zero.

[2] A Parameter Code Reset (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.

[3] A Save Parameters (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 70 and 71). 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 71) 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.

- [4] The Page Control (PC) field defines the type of log parameter the initiator selects to change with the Log Select command. The PC field is defined in Table 69.

Table 69: 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 69.

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).

- [5] The Parameter List Length 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 70 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 note [4] in Table 70). The Page Code (byte 0) specifies which log the page pertains to. The Page Code assignments are listed in Table 73. Initiator cannot send page zero.

- [6] See "Control byte," paragraph 4.2.6.

The Log Sense command and its Data In phase also use the format of Table 70, 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 5.9). This table is not repeated in Section 5.9.

Table 70: Log Page format

Bit Byte	7	6	5	4	3	2	1	0
0	Reserved		Page Code [1]					
1	Reserved							
2	(MSB)							
3	Page Length (n-3) [2]						(LSB)	
Log Parameter Structure(s) [3]								
4 to x+3	Log Parameter (First) [4] (Length X bytes)							
:	: [4]							
n-Y to n	Log Parameter (Last) [4] (Length Y bytes)							

Notes.

- [1] Page Codes implemented are given in Table 73.
- [2] The Page Length field 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 71.

Table 71: Log Parameters

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

Notes.

- [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 73.

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.

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] If Disable Save (DS) is zero, it indicates that the drive supports Log Select and Log Sense data saving for that log parameter. The drive saves the current cumulative and the current threshold parameter values in response to a Log Select or Log Sense command with an SP bit of one. A DS bit of one indicates that the drive 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 typically (may not always be) once every ten minutes. A one bit indicates that the drive does not use its save method.

- [5] Enable Threshold Comparison (ETC). A one indicates that a comparison to the threshold value is performed whenever the cumulative value is updated. A bit of zero indicates the comparison is not performed. The value of the ETC bit is the same for both the threshold and cumulative parameters.
- [6] The Threshold Met Criteria (TMC) field defines the basis for comparison of the cumulative and threshold values. See Table 72 for meanings of values in this field. The TMC field is only valid when the ETC bit is one.

Table 72: 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 96) 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] for this table.

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.

- [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] of this table.

The page code assignments for the log pages are listed in Table 73. Detailed descriptions follow the table.

Table 73: Log Page codes

Page Code	Description	Section	Page
02h	Error Counter page (Write)	5.8.1	107
03h	Error Counter page (Read)	5.8.1	107
05h	Error Counter page (Verify)	5.8.1	107
06h	Non-Medium Error page	5.8.2	108
37h	Cache Statistics page	5.8.3	108
3Eh	Factory Log page	5.8.4	109

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

Table 74 defines the Parameter Code field for the write, read and verify error counter pages.

Table 74: Write, Read, and Verify Errors Parameter Code field (02h, 03h, and 05h)

Bit Byte	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.
- 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.
 - 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.
 - 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.
 - 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.

- 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.
- 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.
- 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.

5.8.2 Non-Medium Error page (06h)

Log page code 06h specifies non-medium errors. The page format is shown in Table 75.

Table 75: Non-Medium Error page (06h)

Bit Byte	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.

5.8.3 Cache Statistics page (37h)

Log Page code 37h specifies Cache Statistics page. The page format is shown in Table 76.

Table 76: Cache Statistics page (37h)

Bit Byte	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 04h are described following:

- 00h This parameter code represents the number of logical blocks that have been sent to an initiator.
- 01h This parameter code represents the number of logical blocks that have been received from an initiator.
- 02h This parameter code represents the number of logical blocks read from the cache memory that have been sent to an initiator.
- 03h This parameter code represents the number of read and write commands that had data lengths equal or less than the current segment size.
- 04h This parameter code represents the number of read and write commands that had data lengths greater than the current segment size.

5.8.4 Factory Log page (3Eh)

Log page code 3Eh specifies factory status parameters.

Table 77: Factory Log page (3Eh)

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

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.
- [2] Parameter Code 08h. This parameter reports the time, in minutes, to the next scheduled interrupt for a S.M.A.R.T. measurement.

5.9 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 Table 78. The format of the data pages sent back by the drive in the Data In phase is shown in Table 70.

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

Table 78: Log Sense command (4Dh)

Bit Byte	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) _____							
6	Parameter Pointer [5]						_____ (LSB)	
7	(MSB) _____							
8	Allocation Length [6]						_____ (LSB)	
9	0	0	0	0	0	0	Flag [7]	Link [7]

Notes.

- [1] A Parameter Pointer Control (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]).
- 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] If the Save Parameters (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 71) to a nonvolatile location, and then performs the rest of the Log Sense command.
- [3] The Page Control (PC) field defines the type of drive log parameter the initiator requests the drive to send back on the Data In phase. See Table 69 for definition of this field. The parameter values returned on the Data In phase are from one of the following:
- 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).
 - 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).
 - 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] The Page Code field specifies the page of data requested by the command (see Table 73). 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 79 (Page Code 00h) returns the list of log pages that the drive supports. This page is not defined for the Log Select command.

Table 79: Supported Log pages

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

Notes.

- [a] The Page Length field specifies the length in bytes of the following supported page list.
- [b] 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.
- [5] The Parameter Pointer 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 (and 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] The Allocation Length 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] See "Control byte," paragraph 4.2.6.

5.10 Mode Select (6) command (15h)

The Mode Select (6) command (Table 80) 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.

Table 80: Mode Select (6) command (15h)

Bit Byte	7	6	5	4	3	2	1	0
0	0	0	0	1	0	1	0	1
1	Logical Unit Number [1]			PF [2]	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 [5]	Link [5]

Notes.

- [1] The Logical Unit Number 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.
- [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.

5.10.1 Mode Select parameter list

The Mode Select parameter list (Table 81) 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 81.

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:

- a. If the Strict mode is enabled (see note [6] of Table 101) 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 com-

mand. (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.)

- b. If the initiator attempts to send an unsupported value or a nonzero value to a reserved field in the Mode Select header, block descriptor, or any page header.
- c. 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.
- d. 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 [5] of Table 101).
- e. If the initiator sends a page descriptor with an unsupported page code value and the Strict mode is enabled (see note [6] of 101).

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:

- a. round the parameter to an acceptable value and, if Round is one, terminate the command as described in Section 4.8, or
- b. 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 81: Mode Select parameter list

Bit Byte	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 or Number of Blocks (MSB) [3] [4]
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) [5]
6	Block Length [5]
7	Block Length (LSB) [5]

Parameter Information [6]

0-n	Mode Select Page Headers and Their Parameters (Tables 82 and 83)
-----	---

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 parame-

ters. 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 may be useable as MSB part of the number in bytes 1, 2, and 3, if the drive's capacity capability requires it. For drives whose programmed capacity or maximum designed capacity does not require the use of byte 0, byte 0 will always be 00h. (See also note [4b]).
- [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 individual drive's Product Manual, Volume 1), a number in bytes 0, 1, 2, and 3 that is less than the maximum number of LBAs changes the drive capacity to the value in the block descriptor bytes 0, 1, 2, and 3. A capacity value greater than the maximum number of LBAs possible is rounded down to the maximum capacity possible for the drive.
- [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 (Section 5.12) for detailed descriptions of the Mode Select/Sense pages.

5.10.2 Mode Select page descriptors

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 (Section 5.12) and are not repeated here.

Table 82: Mode Select page descriptor header

Bit Byte	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.

Note. See individual drive's Product Manual, 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.

Table 83: 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	Notch page
0Dh or 1A	Power Condition page
10h	Xor Control Mode page
19h	SCSI Port Control Mode page
1Ch	Informational 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.

5.11 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. Initiators should issue Mode Sense prior to Mode Select to determine supported pages, page lengths, and other parameters.

Table 84: Mode Select (10) command (55h)

Bit Byte	7	6	5	4	3	2	1	0
0	0	1	0	1	0	1	0	1
1	Logical Unit Number [1]			PF [2]	Reserved			SMP [3]
2 : 6	Reserved							
7	(MSB) _____							
8	Parameter List Length [4]						_____ (LSB)	
9	Control [5]							

Notes.

- [1] The Logical Unit Number 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.
- [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.

5.11.1 Mode Select (10) parameter list

The Mode Select (10) parameter list (Table 85) contains a four byte header, followed by zero or one block descriptor, followed by the pages of Mode Select (10) parameters.

Acceptable values for the Mode Select parameter list for the drive are shown in Table 84.

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:

- a. If the Strict mode is enabled (see note [6] of Table 101) 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.)
- b. If the initiator attempts to send an unsupported value or a nonzero value to a reserved field in the Mode Select header, block descriptor, or any page header.

- c. 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.
- d. 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 [5] of Table 101).
- e. If the initiator sends a page descriptor with an unsupported page code value and the Strict mode is enabled (see note [6] of Table 101).

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:

- a. round the parameter to an acceptable value and, if Round is one, terminate the command as described in Section 4.8, or
- b. 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 85: Mode Select (10) parameter list

Bit Byte	7	6	5	4	3	2	1	0
0,1	0	0	0	0	0	0	0	0
	Reserved							
2	0	0	0	0	0	0	0	0
	Medium Type [1]							
3,4,5	0	0	0	0	0	0	0	0
	Reserved							
6,7	Block Descriptor Length after 0 or 8 (decimal) [2]							
Block Descriptor Data								
0	0	0	0	0	0	0	0	0
	Density Code [3]							
1	0	0	0	0	0	0	0	0
	Number of Blocks (MSB) [3] [4]							
2	0	0	0	0	0	0	0	0
	Number of Blocks [4]							
3	0	0	0	0	0	0	0	0
	Number of Blocks (LSB) [4]							
4	0	0	0	0	0	0	0	0
5	Block Length (MSB) [5]							
6	Block Length [5]							
7	Block Length (LSB) [5]							
Parameter Information [6]								
0-n	Mode Select Page Headers and Their Parameters (Tables 82 and 83)							

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.

ters. 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 individual drive's Product Manual, Volume 1), a number in bytes 0, 1, 2, and 3 that is less than the maximum number of LBAs 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 LBAs 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 (10) command (Section 5.13) for detailed descriptions of the Mode Select/Sense pages.

5.12 Mode Sense (6) command (1Ah)

The Mode Sense (6) 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.

Table 86: Mode Sense (6) command (1Ah)

Bit Byte	7	6	5	4	3	2	1	0
0	0	0	0	1	1	0	1	0
1	Logical Unit Number [1]			0	DBD [2]	0	0	0
2	PCF [3]		Page Code [4]					
3	Block Descriptor Length Field [2]							
4	Allocation Length [5]							
5	0	0	0	0	0	0	Flag [6]	Link [6]

Notes.

- [1] The Logical Unit Number must be zero.
- [2] A Disable Block Descriptors (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. The Block Descriptor Length Field 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 and 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 87.

Table 87: Page Control Field bits affect on mode parameters returned

Bit 7 Bit 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 non-changeable 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's Product Manual, Volume 1):

Page Code	Description
00h	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 Parameters page
0Ah	Control Mode page
0Ch	Notch page
0Dh or 1Ah	Power Condition page
10h	Xor Control Mode page
19h	SCSI Port Control Mode page
1Ch	Informational 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 88: Mode Sense Data

Bit Byte	7	6	5	4	3	2	1	0
0	Sense Data Length [1]							
1	0	0	0	0	0	0	0	0
	Medium Type [2]							
2	WP [3]	0	0	DPO-FUA [8]	0	0	0	0
	Reserved		Reserved		Reserved			
3	0	0	0	0	1	0	0	0
	Block Descriptor Length (8 decimal) [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
	Reserved
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
-----	--

Notes.

- [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 nonzero. See individual drive's Product Manual, Volume 1, 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 DPO-FUA bit of zero indicates that the target does not contain a cache memory or does not support the DPO and FUA bits. A DPO-FUA bit of one indicates that the target supports the DPO and FUA bits.

5.12.1 Mode Sense page descriptor header

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 individual drive's Product Manual, Volume 1, for a table giving the Mode Sense parameter values that are applicable to the drive model of interest. The tables in 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.

Table 89: Mode Sense page descriptor header

Bit	7	6	5	4	3	2	1	0
Byte								
0	PS	0	Page Code [1]					
1	Page Length [1]							
2–n	Mode Parameters [1]							

Notes.

[1]	Page Code	Page Description	Reference Table	Page
	01h	Error Recovery	90	123
	02h	Disconnect/Reconnect Control	91	126
	03h	Format Parameters	92	128
	04h	Rigid Drive Geometry	93	130
	07h	Verify Error Recovery	94	132
	08h	Caching Parameters	95	134
	0Ah	Control Mode	96/97	137/138
	0Ch	Notch	98	139
	0Dh or 1Ah	Power Condition	99	141
	10h	Xor Control Mode	100	142
	00 h	Unit Attention Parameters	101	144
	1Ch	Informational Exceptions Control	102	146

5.12.2 Error Recovery page (01h)

The drive Error Recovery page implementation is defined in Table 90. This table summarizes the function, the default value, and changeability status for each byte/bit.

Table 90: Error Recovery page (01h)

Bit	7	6	5	4	3	2	1	0	
Byte									
Page Descriptor Header									
0	1	0	0	0	0	0	0	1	
	PS [1]		Page Code (01h)						
1	Page Length (0Ah)								
Error Recovery Parameters									
2 Default	AWRE [3]	ARRE [4]	TB [5]	RC [6]	EER [7]	PER [8]	DTE [9]	DCR [10]	
Changeable	[2]								
3 Default	Read Retry Count [11]								
Changeable	[2]								
4 Default	Correction Span (bits) [12]								
Changeable	[2]								
5 Default	Head Offset Count [13]								
Changeable	[2]								
6 Default	Data Strobe Offset Count [14]								
Changeable	[2]								
7 Default	Reserved [17]								
Changeable	[2]								
8 Default	Write Retry Count [15]								
Changeable	[2]								
9 Default	Reserved								
Changeable	[2]								
10	(MSB)	Recovery Time Limit [16]							
11								(LSB)	
Changeable	[2]								

Notes.

- [1] The returned Parameter Savable (PS) 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 individual drive's Product Manual, 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.
- [4] 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.
- [5] 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.
- [6] 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 individual drive's Product Manual, Volume 1, section showing changeable values.
- [7] 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.
- [8] 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.
- [9] 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.
- [10] The Disable Correction (DCR) bit, when set to one, indicates that only on-the-fly ECC correction shall be applied. The exact behavior is product specific. When set to zero, maximum ECC correction shall be applied if correction is possible.
- [11] 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 individual drive's Product Manual, Volume 1, for more details on the levels of error recovery available.
- [12] 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.
- [13] 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.
- [14] 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.
- [15] 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 individual drive's Product Manual, Volume 1, for more details on the levels of error recovery available.

- [16] 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 or 0000h means that the command recovery time is unlimited. A value of 0001h 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 drive's Product Manual, Volume 1, Mode page changeable bit settings for Mode page 01h, bytes 10 and 11.
- [17] In SCSI-1 mode of operation, this byte is the recovery time limit value (see note [16]).

5.12.3 Disconnect/Reconnect Control page (02h)

The Disconnect/Reconnect Page implementation is defined in Table 91. This table summarizes the function and defines the default value and changeable status.

Table 91: Disconnect/Reconnect Control page (02h)

Bit	7	6	5	4	3	2	1	0	
Byte									
Page Descriptor Header									
0	1	0	0	0	0	0	1	0	
	PS [1]		Page Code (02h)						
1	Page Length (0Eh)								
Disconnect/Reconnect Control Parameters									
2 Default	Buffer Full Ratio [3] [4]								
Changeable	[2]								
3 Default	Buffer Empty Ratio [3] [5]								
Changeable	[2]								
4 Default	Bus Inactivity Limit (MSB) [6]								
Changeable	[2]								
5 Default	Bus Inactivity Limit (LSB) [6]								
Changeable	[2]								
6,7 Default	Disconnect Time Limit [7]								
Changeable	[2]								
8,9 Default	Connect Time Limit [8]								
Changeable	[2]								
10,11 Default	(MSB)	Maximum Burst Size [9]						(LSB)	
Changeable	[2]								
12 Default	EMDP [11]	0	0	0	DImm [12]	DTDC [10]			
Changeable	[2]								
13-15	Reserved								

Notes.

- [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 individual drive's Product Manual, Volume 1, 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] The Buffer Full Ratio indicates, on Read commands, how full the drive's buffer shall be before attempting a reselection. This value is changeable by an initiator.

- [5] The Buffer Empty Ratio indicates, on Write commands, how empty the drive's buffer shall be before attempting a reselection. This value is changeable by an initiator, but the value is ignored by the drive. The drive uses an adaptive algorithm that accounts for the buffer segment size, the command transfer length, and data transfer rate to determine the optimal point to reconnect.
- [6] 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.
- [7] 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 drive's Product Manual, Volume 1).
- [8] 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.
- [9] 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.
- [10] 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 individual drive's Product Manual, Volume 1, which indicates if a particular drive implements the DTDC function.

- [11] 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.
- [12] 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 drive's Product Manual, Volume 1.

5.12.4 Format Parameters page (03h)

The Format Parameters page implementation is defined in Table 92. This table summarizes the function and defines the default or changeability status for each bit. The actual implementation of reserving spare areas for defect management takes place during the Format Unit command.

Note. In Table 92, zone refers to 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).

Table 92: Format Parameters page (03h)

Bit	7	6	5	4	3	2	1	0
Byte								
Page Descriptor Header								
0	1	0	0	0	0	0	1	1
	PS [1]		Page Code (03h)					
1	Page Length (16h)							
Format Parameters [1]								
2,3 Default	Tracks per Zone (MSB) [2]							
Changeable	[12]							
4,5 Default	Alternate Sectors per Zone [3]							
Changeable	[12]							
6,7 Default	Alternate Tracks per Zone [4]							
Changeable	[12]							
8,9 Default	Alternate Tracks per Volume [5]							
Changeable	[12]							
10,11 Default	Sectors per Track [6]							
Changeable	[12]							
12,13 Default	Data Bytes per Physical Sector [7]							
Changeable	[12]							
14,15 Default	Interleave [8]							
Changeable	[12]							
16,17 Default	Track Skew Factor [9]							
Changeable	[12]							
18,19 Default	Cylinder Skew Factor [10]							
Changeable	[12]							
20 Default	SSEC [11]	HSEC [11]	RMB [11]	SURF [11]	Reserved [11]			
Changeable	[12]							
21,22,23 Default	Reserved							
Changeable	[12]							

Notes.

- [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's Product Manual, Volume 1, 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 Data 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.
- [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. Bits 0-5, and 7 are not implemented by the drive and are always zero. All bits (0-7) are not changeable.
- [12] See individual drive's Product Manual, Volume 1, Mode Sense Data section for changeable values.

5.12.5 Rigid Drive Geometry Parameters page (04h)

The Rigid Drive Geometry Parameters page implementation is defined in Table 93. This table summarizes the function and defines the default value

Table 93. Rigid Drive Geometry Parameters page (04h)

Bit	7	6	5	4	3	2	1	0	
Byte									
Page Descriptor Header									
0	1	0	0	0	0	1	0	0	
	PS [1]		Page Code (04h)						
1	Page Length (16h)								
Rigid Drive Geometry Parameters									
2 Default	Number of Cylinders (MSB) [2]								
Changeable	[8]								
3 Default	Number of Cylinders [2]								
Changeable	[8]								
4 Default	Number of Cylinders (LSB) [2]								
Changeable	[8]								
5 Default	Number of Heads [3]								
Changeable	[8]								
6,7,8 Default	Starting Cylinder—Write Precomp [4]								
Changeable	[8]								
9,10,11 Default	Starting Cylinder—Reduced Write Current [4]								
Changeable	[8]								
12,13 Default	Drive Step Rate [4]								
Changeable	[8]								
14,15,16 Default	Landing Zone Cylinder [4]								
Changeable	[8]								
17 Default	0	0	0	0	0	0	RPL [5]		
Changeable	[8]								
18 Default	Rotational Offset XXh [6]								
Changeable	[8]								
19 Default	0	0	0	0	0	0	0	0	
	Reserved								
Changeable	[8]								

Table 93. Rigid Drive Geometry Parameters page (04h) (continued)

Bit Byte	7	6	5	4	3	2	1	0
20	(MSB)							
21	Medium Rotation Rate [7]							(LSB)
Changeable	[8]							
22 Default	0	0	0	0	0	0	0	0
Changeable	Reserved							
23 Default	0	0	0	0	0	0	0	0
Changeable	Reserved							
Changeable	[8]							

Notes.

- [1] The Parameter Savable (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 80) is 1. See individual drive's 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's 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 (rotational 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 individual drive's Product Manual, Volume 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 individual drive's Product Manual, Volume 1, Mode Sense Data section, for changeable values.

5.12.6 Verify Error Recovery page (07h)

The Verify Error Recovery page (Table 94) specifies the error recovery parameters the target shall use during the Verify command and the verify operation of the Write and Verify command.

Table 94: Verify Error Recovery page (07h)

Bit Byte	7	6	5	4	3	2	1	0
Page Descriptor Header								
0	1 PS [1]	0 Reserved	Page Code (07h)					
1	Page Length (0Ah)							
Verify Error Recovery Parameters								
2 Default	0	0	0	0	EER [2]	PER [3]	DTE [4]	DCR [5]
Changeable	[9] Reserved							
3 Default	Verify Retry Count [6]							
Changeable	[9]							
4 Default	Verify Correction Span (bits) [7]							
Changeable	[9]							
5 Default	0	0	0	0	0	0	0	0
Changeable	[9] Reserved							
6 Default	0	0	0	0	0	0	0	0
Changeable	[9] Reserved							
7 Default	0	0	0	0	0	0	0	0
Changeable	[9] Reserved							
8 Default	0	0	0	0	0	0	0	0
Changeable	[9] Reserved							
9 Default	0	0	0	0	0	0	0	0
Changeable	[9] Reserved							
10 Default	1 (MSB)	1	1	1	1	1	1	1
Changeable	[9] Verify Recovery Time Limit [8]							
11 Default	1	1	1	1	1	1	1	1 (LSB)
Changeable	[9] Verify Recovery Time Limit [8]							

Notes.

- [1] The Parameter 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 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.
- [3] 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.
- [4] 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.
- [5] The Disable Correction (DCR) bit, when set to one, indicates that only on-the-fly ECC correction shall be applied. The exact behavior is product specific. When set to zero, maximum ECC correction shall be applied if correction is possible.
- [6] 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 individual drive's Product Manual, Volume 1, for more details on the levels of error recovery available.
- [7] 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.
- [8] The Verify 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 drive's Product Manual, Volume 1, 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 EER bit to zero, the PER, DTE, and DCR bits to one, and the number of retries and recovery time limit to zero.
- [9] See individual drive's Product Manual, Volume 1, Mode Sense Data section, for changeable values.

5.12.7 Caching Parameters page (08h)

The Caching Parameters page for Mode Sense/Mode Select defines the parameters that affect the use of the cache (Table 95).

Table 95: Caching Parameters page (08h)

Bit Byte	7	6	5	4	3	2	1	0
Page Descriptor Header								
0	PS [1]	Reserved	Page Code (08h)					
1	Page Length (12h)							
2	IC [3]	ABPF [4]	CAP [5]	DISC [6]	SIZE [7]	WCE [8]	MF [9]	RCD [10]
Changeable	[2]							
3	Demand Read Retention Priority [11]				Write Retention Priority [12]			
Changeable	[2]							
4	(MSB)	Disable Prefetch Transfer Length [13]						(LSB)
5								
Changeable	[2]							
6	(MSB)	Minimum Prefetch [14]						(LSB)
7								
Changeable	[2]							
8	(MSB)	Maximum Prefetch [15]						(LSB)
9								
Changeable	[2]							
10	(MSB)	Maximum Prefetch Ceiling [16]						(LSB)
11								
Changeable	[2]							
12	FSW [17]	LBCSS [18]	DRA [19]	Reserved				
Changeable	[2]							
13	Number of Cache Segments [20]							
Changeable	[2]							
14	(MSB)	Cache Segment Size [21]						(LSB)
15								
Changeable	[2]							
16	Reserved							
17	(MSB)	Non-Cache Segment Size [22]						(LSB)
18								
19								
Changeable	[2]							

Notes.

- [1] The returned Parameter Savable (PS) 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 individual drive's Product Manual, Volume 1, section showing changeable values.

- [3] When the Initiator Control (IC) enable bit is 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 ARLA (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's Product Manual, Volume 1.
- [4] The Abort Prefetch (ABPF) bit, when set to one, with the DRA bit equal to zero, requests that the SCSI device abort the prefetch upon selection. The ABPF set to one takes precedence over the Minimum Prefetch bytes. When set to zero, with the DRA bit equal to zero, the termination of any active prefetch is dependent upon Caching Page bytes 4 through 11 and is operation and/or vendor specific.
- [5] The Caching Analysis Permitted (CAP) bit, when set to one, enables caching analysis. A zero indicates caching analysis is disabled. Caching analysis results are placed in the SCSI logging information table, (see Table 76). See individual drive's Product Manual, Volume 1, SCSI Bus Conditions and Miscellaneous Features Supported table.
- [6] The Discontinuity (DISC) bit, when set to one, requests that the SCSI device continue the prefetch across time discontinuities, such as across cylinders or tracks up to the limits of the buffer, or segment, space available for prefetch. When set to zero, the DISC requests that prefetches be truncated at time discontinuities.
- [7] The Size Enable (SIZE) bit, 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.
- [8] Write Cache Enable (WCE).
- 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.
- [9] Multiplication Factor (MF).
- 0: The Minimum Prefetch and Maximum Prefetch fields are interpreted as a number of logical blocks.
 - 1: Specifies that the target shall interpret the minimum and maximum prefetch 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 prefetch.
- [10] Read Cache Disable (RCD).
- 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.
- [11] 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.
- [12] 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.
- [13] Disable Prefetch Transfer Length. Prefetch is disabled for any SCSI Read command whose requested transfer length exceeds this value.
- [14] Minimum Prefetch specifies the minimum number sectors to prefetch, regardless of the delay it may cause to other commands.
- [15] Maximum Prefetch specifies the maximum number of logical blocks that may be prefetched. The prefetch operation may be aborted before the maximum prefetch value is reached, but only if the minimum prefetch value has been satisfied.
- [16] Maximum Prefetch Ceiling specifies an upper limit on the number of logical blocks computed as the maximum prefetch. If the Maximum Prefetch value is greater than the Maximum Prefetch Ceiling, the value is truncated to the Maximum Prefetch Ceiling value.

- [17] The Force Sequential Write (FSW) bit, 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.
- [18] The Logical Block Cache Segment Size (LBCSS) bit is not used at this time.
- [19] The Disable Read-Ahead (DRA) bit, 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).
- [20] The Number of Cache Segments byte gives the number of segments into which the host requests the drive divide the cache.
- [21] The Cache Segment Size field 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.
- [22] The Non-Cache Segment Size field, when greater than zero, 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 Segment Size equal 0 or the same of this field plus the Cache Segment Size greater than the buffer size is vendor specific.

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.

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.

5.12.8 Control Mode page (0Ah)

The Mode Sense/Mode Select Control Mode page (Table 96) 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.

Table 96: Control Mode page (0Ah)

Bit Byte	7	6	5	4	3	2	1	0
0	1 PS [1]	Reserved	Page Code (0Ah)					
1	Page Length (0Ah)							
2 Default	0	0	0	0	0	0	GLTSD [9]	RLEC [2]
Changeable	[7]							
3 Default	Queue Algorithm Modifier [3]				0	0	QErr [4]	DQue [5]
Changeable	[7]							
4 Default	EECA [6]	RAC [10]	0	0	0	0	0	0
Changeable	[7]							
5 Default	0	0	0	0	0	0	0	0
Changeable	[7]							
6 & 7 Default	0	0	0	0	0	0	0	0
Changeable	[7]							
8 & 9 Default	Busy Timeout Period [8]							
Changeable	[7]							
10	Reserved (all zeros)							
11	Reserved (all zeros)							

Notes.

- [1] The Parameter Savable (PS) 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 97: Queue Algorithm Modifier

Value	Definition
0h	Guaranteed data integrity: 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.
1h	Unrestricted re-ordering allowed: 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.
2h - 7h	Reserved.
8h - Fh	Vendor Specific.

- [4] A Queue Error Management (QEerr) 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 QEerr 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 individual drive's Product Manual, Volume 1, 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] A Global Logging Target Save Disable (GLTSD) 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] Report A Check (RAC). Not used at this time.

5.12.9 Notch page (0Ch)

The Notch page (Table 98) 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 98: Notch page (0Ch)

Bit Byte	7	6	5	4	3	2	1	0
0	PS [1]	0 Reserved	0	0	1	1	0	0
	Page Code (0Ch)							
1	Page Length (16h)							
2 Default	ND [2]	LPN [3]	0	0	0	0	0	0
	Reserved							
Changeable	[9]							
3 Default	0	0	0	0	0	0	0	0
	Reserved							
Changeable	[9]							
4,5 Default	(MSB)	Maximum Number of Notches [4]						(LSB)
Changeable	[9]							
6,7 Default	(MSB)	Active Notch [5]						(LSB)
Changeable	[9]							
8-11 Default	(MSB)	Starting Boundary [6]						(LSB)
Changeable	[9]							
12-15 Default	(MSB)	Ending Boundary [7]						(LSB)
Changeable	[9]							
16-23 Default	(MSB)	Pages Notched [8]						(LSB)
Changeable	[9]							

Notes.

- [1] The Parameter 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. An 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. An 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.
- [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. An 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's Product Manual, Volume 1, 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.

5.12.10 Power Condition page (0Dh or 1Ah [5])

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.

Table 99: Power Condition page (0Dh or 1Ah [5])

Bit Byte	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]							

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 individual drive's Product Manual, Volume 1, Mode Sense Data and Commands Supported tables for indication of which the drive model in question uses.
- [6] See individual drive's Product Manual, Volume 1, 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.

5.12.11 Xor Control Mode page (10h)

The Xor Control Mode page (Table 100) provides the initiator with the means to obtain or modify certain Xor operating parameters of the target.

Table 100: Xor Control Mode page (10h)

Bit Byte	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)							
5	Maximum Xor Write Size [2]							
6								
7								
Changeable	[6]							
8	(MSB)							
9	Reserved							
10								
11								
Changeable	[6]							
12	(MSB)							
13	Maximum Regenerate Size [3]							
14								
15								
Changeable	[6]							
16	(MSB)							
17	Maximum Rebuild Read Size [4]							
18								
19								
Changeable	[6]							
20	Reserved							
21								
Changeable	[6]							
22	Rebuild Delay [5]							
23								
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's Product Manual, Volume 1, 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.

5.12.12 Unit Attention Parameters page (00h)

The Unit Attention Parameters page is the last page to be reported by the drive.

Table 101: Unit Attention Parameters page (00h)

Bit	7	6	5	4	3	2	1	0
0	1	0	0	0	0	0	0	0
	PS [10]		Page Code (00h)					
1	Page Length (in bytes) [1]							
2 Default	0	SSM [2]	Inquiry Length [3]	Unit Attention [4]	0	Round [5]	Strict [6]	SCSI-2 [7]
Changeable	[9]							
3 Default	Reserved [8]	Self Seek [11]	Reserved [8]					
Changeable	[9]							

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 message, 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] When the Inquiry Length 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 Attention 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, and 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.

-
- [7] When set to one, the SCSI-2 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 (08h) Length from 12h to 0Ah.
- [8] These bits are reserved for future compatibility with Seagate host adapters. Though they presently may be changeable, (see note [7]) these bits do not control anything, unless the individual drive's Product Manual, Volume 1, indicates that they do and defines their use in the Mode Sense Data section.
- [9] See individual drive's Product Manual, Volume 1, 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 Parameter Savable (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).
- [11] If the Self Seek bit is set to one, the drive will enter self seek mode for testing purposes. Such testing could include, but is not limited to, power dissipation and acoustics. While in this mode, the drive will accept SCSI commands and will process them in between the self seek operations, including a mode select to turn this bit back off. As such, this bit should be off for normal drive operations. If this bit is set to zero, the drive will not self seek; normal operating mode.

5.12.13 Informational Exceptions Control page (1Ch)

The Informational Exceptions Control page (Table 102) 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 102: Informational Exceptions Control page (1Ch)

Bit	7	6	5	4	3	2	1	0	
Byte									
Page Descriptor Header									
0	PS	Reserved	Page Code (1Ch)						
1	Page Length (0Ah)								
Informational Exceptions Control Parameters									
2	Perf [1]	Reserved	EWasc [3]	DExcpt [2]	Test [4]	Reserved	LogErr [5]		
Changeable	[10]								
3	Reserved				MRIE [6]				
Changeable	[10]								
4	Interval Timer [7] [9]								
5									(MSB)
6									
7									(LSB)
Changeable	[10]								
8	Report Count [8] [9]								
9									(MSB)
10									
11									(LSB)
Changeable	[10]								

Notes.

- [1] A Performance (Perf) bit 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 the failure prediction threshold exceeded reporting shall be enabled. The method of reporting the failure prediction threshold exceeded 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 reporting of the failure prediction threshold exceeded. The Method of Reporting Informational Exceptions field is ignored when DExcpt is set to one and EWasc is set to zero.
- [3] An Enable Warning (EWasc) bit of zero indicates the target shall disable reporting of the warning. The Method of Reporting Informational Exceptions field is ignored when DExcpt is set to one and EWasc is set to zero. An EWasc bit of one indicates warning reporting shall be enabled. The method for reporting the warning when the EWasc bit is set to one is determined from the Method of Reporting Information Exceptions field.

- [4] A Test bit 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.
- [5] The Log Errors (LogErr) bit 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.
- [6] The Method of Reporting Informational Exceptions (MRIE) field indicates the methods that shall be used by the target to report informational exception conditions (see Table 103). The priority of reporting multiple information exceptions is vendor specific.
- [7] 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.
- [8] 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.
- [9] 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.]
- [10] See individual drive's Product Manual, Volume 1, 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 103: Format of Method of Reporting Informational Exceptions (MRIE) field

Code	Description
0h	No reporting of informational exception condition. This method instructs the target to not report information exception conditions.
1h	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. The sense key shall be set to Recovered Error and the additional sense code shall indicate the cause of the informational exception condition.
2h	Generate Unit Attention. This method instructs the target to report informational exception conditions by returning a Check Condition status on any command. The sense key shall be set to Unit Attention and the additional sense code shall indicate the cause of the informational exception condition. The command that has the Check Condition shall not be executed before the informational exception condition is reported.
3h	Conditionally generate recovered error. This method instructs the target to report informational exception conditions, dependent on the value of the PER bit of the error recovery parameters mode page, by returning a Check Condition status on any command. The sense key shall be set to Recovered Error and the additional sense code shall indicate the cause of the informational exception condition. The command that has the Check Condition shall complete without error before any informational exception condition may be reported.
4h	Unconditionally generate recovered error. This method instructs the target to report informational exception conditions, regardless of the value of the PER bit of the error recovery parameters mode page, by returning a Check Condition status on any command. The sense key shall be set to Recovered Error and the additional sense code shall indicate the cause of the informational exception condition. The command that has the Check Condition shall complete without error before any informational exception condition may be reported.
5h	Generate no sense. This method instructs the target to report informational exception conditions by returning a Check Condition status on any command. The sense key shall be set to No Sense and the additional sense code shall indicate the cause of the informational exception condition. The command that has the Check Condition shall complete without error before any informational exception condition may be reported.
6h	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.
7h - Bh	Reserved.
Ch - Fh	Vendor Specific.

5.13 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.

Table 104: Mode Sense (10) command (5Ah)

Bit Byte	7	6	5	4	3	2	1	0
0	0	1	0	1	1	0	1	0
1	Logical Unit Number [1]			Reserved	DBD [2]	Reserved		
2	PCF [3]		Page Code [4]					
3 : 6	Reserved							
7	(MSB) _____							
8	Allocation Length [5]						_____ (LSB)	
9	Control [6]							

Notes.

- [1] The Logical Unit Number must be zero.
- [2] A Disable Block Descriptors (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 and 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 105.

Table 105: Page Control Field bits affect on mode parameters returned

Bit 7	Bit 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 non-changeable 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's Product Manual, Volume 1):

Page Code	Description
00h	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 Parameters page
0Ah	Control Mode page
0Ch	Notch page
0Dh or 1Ah	Power Condition page (see individual drive's Product Manual, Volume 1, for code used)
10h	Xor Control Mode page
1Ch	Informational 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 106: Mode Sense Parameter List

Bit Byte	7	6	5	4	3	2	1	0
0,1	Sense Data Length [1]							
2	0	0	0	0	0	0	0	0
Medium Type [2]								
3	WP [3]	0	0	DPO-FUA [8]	0	0	0	0
		Reserved				Reserved		
4,5	Reserved							
6,7	0	0	0	0	1	0	0	0
Block Descriptor Length (8 decimal) [4]								
Block Descriptor Data								
0	0	0	0	0	0	0	0	0
Density Code [5] [6]								
1	0	0	0	0	0	0	0	0
Number of Blocks (MSB) [6]								
2	0	0	0	0	0	0	0	0
Number of Blocks [6]								
3	0	0	0	0	0	0	0	0
Number of Blocks (LSB) [6]								
4	0	0	0	0	0	0	0	0
Reserved								
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							

Notes.

- [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 nonzero. See individual drive's Product Manual, Volume 1, 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 DPO-FUA bit of zero indicates that the target does not contain a cache memory or does not support the DPO and FUA bits. A DPO-FUA bit of one indicates that the target supports the DPO and FUA bits.

5.14 Move Medium command (A7h)

Not implemented. If received, a Check Condition status is sent.

5.15 Persistent Reserve In command (5Eh)

The Persistent Reserve In command (see Table 107) is used to obtain information about persistent reservations and reservation keys that are active within a device server. This command is used in conjunction with the Persistent Reserve Out command (see Section 5.16). For additional information about Reservations, see Section 5.35.5.

When a device server receives a Persistent Reserve In command during the same time Reserve (10) or Reserve (6) logical unit or extent reservations or SMC element reservations are active (see 5.35 and 5.36), the command shall be rejected with a Reservation Conflict status. A Persistent Reserve In command shall not conflict with any persistent reservation.

Table 107: Persistent Reserve In command (5Eh)

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

Notes.

- [1] The actual length of the Persistent Reserve In parameter data is available in a parameter data field. The Allocation Length field in the CDB indicates how much space has been reserved for the returned parameter list. If the length is not sufficient to contain the entire parameter list, the first portion of the list shall be returned. This shall not be considered an error. If the remainder of the list is required, the application client should send a new Persistent Reserve In command with an Allocation Length field large enough to contain the entire list.
- [2] The Service action codes for the Persistent Reserve In command are defined in Table 108.

Table 108: Persistent Reserve In Service Action codes

Code	Name	Description
00h	Read Keys	Reads all registered Reservation Keys. The Read Keys service action requests that the device server return a parameter list containing a header and a complete list of all reservation keys currently registered with the device server. If multiple initiators have registered with the same key, then that key value shall be listed multiple times, once for each such registration. The keys may have been passed by a Persistent Reserve Out command that has performed a Register service action. The relationship between a reservation key and the initiator or port is outside the scope of this manual.
01h	Read Reservations	Reads all current persistent reservations. The Read Reservations service action requests that the device server return a parameter list containing a header and a complete list of all unique persistent reservations that are presently active in the device server. Multiple persistent reservations may be returned only if element reservations are present.
02h-1Fh	Reserved	Reserved.

[3] See “Control byte,” paragraph 4.2.6.

5.15.1 Persistent Reserve In parameter data for Read Keys

The format for the parameter data provided in response to a Persistent Reserve In command with the Read Keys service action is shown in Table 170.

Table 109: Persistent Reserve In parameter data for Read Keys

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB) _____							
3	Generation [1]							(LSB)
4	(MSB) _____							
7	Additional Length (n-7) [2]							(LSB)
Reservation Key List [3]								
8	(MSB) _____							
15	First Reservation Key							(LSB)
	.							
	.							
	.							
n-7	(MSB) _____							
n	Last Reservation Key							(LSB)

Notes.

- [1] The Generation value is a 32-bit counter in the device server that shall be incremented every time a Persistent Reserve Out command requests a Register, a Clear, a Preempt, or a Preempt and Abort service action. The counter shall not be incremented by a Persistent Reserve In command, by a Persistent Reserve Out command that performs a Reserve or Release service action, or by a Persistent Reserve Out command that is not performed due to an error or reservation conflict. Regardless of the APTPL bit value, the Generation value shall be set to 0 as part of the power on reset process.

The Generation value allows the application client examining the generation value to verify that the configuration of the initiators attached to a logical unit has not been modified by another application client without the knowledge of the examining application client.

- [2] The Additional Length field contains a count of the number of bytes in the Reservation Key List. If the Allocation Length specified by the Persistent Reserve In command is not sufficient to contain the entire parameter list, then only the bytes from 0 to the maximum allowed Allocation Length shall be sent to the application client. The remaining bytes shall be truncated, although the Additional Length field shall still contain the actual number of bytes in the Reservation Key List without consideration of any truncation resulting from an insufficient Allocation Length. This shall not be considered an error.
- [3] The Reservation Key List contains all the 8-byte reservation keys registered with the device server through Persistent Reserve Out Reserve, Preempt, Preempt and Clear, or Register service actions. Each reservation key may be examined by the application client and correlated with a set of initiators and SCSI ports by mechanisms outside the scope of this manual.

5.15.2 Persistent Reserve In parameter data for Read Reservations

The format for the parameter data provided in response to a Persistent Reserve In command with the Read Reservations service action is shown in Table 171.

Table 110: Persistent Reserve In parameter data for Read Reservations

Bit Byte	7	6	5	4	3	2	1	0
0	Generation [1]							
3								
4	Additional Length (n-7) [2]							
7								
8	Reservations Descriptor [3]							
n								

Notes.

- [1] The Generation value is a 32-bit counter in the device server that shall be incremented every time a Persistent Reserve Out command requests a Register, a Clear, a Preempt, or a Preempt and Clear operation. The counter shall not be incremented by a Persistent Reserve In command, by a Persistent Reserve Out command that performs a Reserve or Release service action, or by a Persistent Reserve Out command that is not performed due to an error or reservation conflict. The Generation value shall be set to 0 as part of the power on reset process.

The Generation value allows the application client examining the generation value to verify that the configuration of the initiators attached to a logical unit has not been modified by another application client without the knowledge of the examining application client.
- [2] The Additional Length field contains a count of the number of bytes in of Reservations Descriptor. If the Allocation Length specified by the Persistent Reserve In command is not sufficient to contain the entire parameter list, then only the bytes from 0 to the maximum allowed Allocation Length shall be sent to the application client. The remaining bytes shall be truncated, although the Additional Length field shall still contain the actual number of bytes of Reservations Descriptor and shall not be affected by the truncation. This shall not be considered an error.
- [3] The format of a single read Reservations Descriptor is defined in Table 111. There shall be a Reservations Descriptor for the persistent reservation, if any, present in the logical unit and a Reservations Descriptor for each element, if any, having a persistent reservation.

Table 111: Persistent Reserve In Read Reservations Descriptor

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB) _____							
7	Reservation Key [1]							(LSB)
8	(MSB) _____							
11	Scope-specific Address [2]							(LSB)
12	Reserved							
13	Scope [3]				Type [4]			
14	_____							
15	Obsolete [2]							_____

Notes.

- [1] If a persistent reservation is present in the logical unit that does contain elements, there shall be a Reservation descriptor in the list of parameter data returned by the device server in response to the Persistent Reserve In command with a Read Reservation service action for the Logical Unit persistent reservation that is held, if any, and each element persistent reservation that may be held. The Reservation descriptor shall contain the Reservation Key under which the persistent reservation is held. The Type and Scope of the persistent reservation as present in the Persistent Reserve Out command that created the persistent reservation shall be returned [3] [4].
- [2] If the Scope is an Element reservation, the Scope-specific Address field shall contain the Element address, zero filled in the most significant bytes to fit the field. If the Scope is a Logical Unit reservation, the Scope-specific Address fields shall be set to zero. The Obsolete field in Bytes 14 and 15 was used for extent reservations as described in an earlier standard.
- [3] The value in the Scope field shall indicate whether a persistent reservation applies to an entire logical unit, to a portion of the logical unit defined as an extent, or to an element. The values in the Scope field are defined in Table 112.

Table 112: Persistent Reservation Scope Codes

Code	Name	Description
00h	LU	A Scope field value of LU shall indicate that the persistent reservation applies to the entire logical unit. The LU scope shall be implemented by all device servers that implement Persistent Reserve Out.
01h	Extent	Obsolete.
02h	Element	Persistent reservation applies to the specified element. A Scope field value of Element shall indicate that the persistent reservation applies to the element of the logical unit defined by the Scope-specific address field in the Persistent Reserve Out parameter list. An element is defined by the SCSI-3 Medium Changer Commands (SMC) standard. The Element scope is optional for all device servers that implement Persistent Reserve Out.
03h-1Fh	Reserved	Reserved.

- [4] The value in the Type field shall specify the characteristics of the persistent reservation being established within the element or within the logical unit. Table 113 defines the characteristics of the five different type values. For each persistent reservation type, Table 113 lists code value and describes the required device server support. In Table 113, the description of required device server support is divided into three paragraphs. The first paragraph defines the required handling for read operations. The second paragraph defines the required handling for write operations.

Table 113: Persistent Reservation Type Codes

Code	Name	Description
00h	Read Shared	<p>Reads Shared: Any application client on any initiator may execute commands that perform transfers from the storage medium or cache of the logical unit to the initiator.</p> <p>Writes Prohibited: Any command from any initiator that performs a transfer from the initiator to the storage medium or cache of the logical unit shall result in a reservation conflict.</p>
01h	Write Exclusive	<p>Reads Shared: Any application client on any initiator may execute commands that perform transfers from the storage medium or cache of the logical unit to the initiator.</p> <p>Writes Exclusive: Any command from any initiator other than the initiator holding the persistent reservation that performs a transfer from the initiator to the storage medium or cache of the logical unit shall result in a reservation conflict.</p>
02h	Read Exclusive	<p>Reads Exclusive: Any command from any initiator other than the initiator holding the persistent reservation that performs a transfer from the storage medium or cache of the logical unit to the initiator shall result in a reservation conflict.</p> <p>Writes Shared: Any application client on any initiator may execute commands that perform transfers from the initiator to the storage medium or cache of the logical unit.</p>
03h	Exclusive Access	<p>Reads Exclusive: Any command from any initiator other than the initiator holding the persistent reservation that performs a transfer from the storage medium or cache of the logical unit to the initiator shall result in a reservation conflict.</p> <p>Writes Exclusive: Any command from any initiator other than the initiator holding the persistent reservation that performs a transfer from the initiator to the storage medium or cache of the logical unit shall result in a reservation conflict.</p>
04h	Shared Access	<p>Reads Shared: Any application client on any initiator may execute commands that perform transfers from the storage medium or cache of the logical unit to the initiator.</p> <p>Writes Shared: Any application client on any initiator may execute commands that perform transfers from the initiator to the storage medium or cache of the logical unit.</p>
05h	Write Exclusive, Registrants Only	<p>Reads Shared: Any application client on any initiator may execute commands that perform transfers from the storage medium or cache of the logical unit to the initiator.</p> <p>Writes Exclusive: Any command from an initiator that has not previously performed a Register service action with the device server that performs a transfer to the storage medium or cache of the logical unit, shall result in a reservation conflict.</p>
06h	Exclusive Access, Registrants Only	<p>Reads Exclusive: Any command from an initiator that has not previously performed a Register service action with the device server that performs a transfer from the storage medium or cache of the logical unit, shall result in a reservation conflict.</p> <p>Writes Exclusive: Any command from an initiator that has not previously performed a Register service action with the device server that performs a transfer to the storage medium or cache of the logical unit, shall result in a reservation conflict.</p>
07h-1Fh	Reserved	

5.16 Persistent Reserve Out command (5Fh)

The Persistent Reserve Out command (see Table 114) is used to reserve a logical unit for the exclusive or shared use of a particular initiator. The command shall be used in conjunction with the Persistent Reserve In command and shall not be used with the Reserve and Release commands. For additional information about Reservations, see Section 5.35.5.

Initiators performing Persistent Reserve Out Service actions are identified by a reservation key provided by the application client. An application client may use the Persistent Reserve In command to identify which initiators are holding a persistent reservation and use the Persistent Reserve Out command to preempt that reservation if required.

If a Persistent Reserve Out command is attempted, but there are insufficient device server resources to complete the operation, the device server shall return a Check Condition status. The sense key shall be set to Illegal Request and the additional sense data shall be set to Insufficient Registration Resources.

Table 114: Persistent Reserve Out command (5Fh)

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

Notes.

[1] The Persistent Reserve Out command contains fields that specify a persistent reservation Service Action, the intended scope of the persistent reservation, and the restrictions caused by the persistent reservation. The Type and Scope fields are defined in notes [3] and [4] of Table 111. If a Scope field specifies a scope that is not implemented, the device server shall return a Check Condition status. The sense key shall be set to Illegal Request and additional sense data shall be set to Invalid Field in CDB.

Fields contained in the Persistent Reserve Out parameter list specify the reservation keys and extent information required to perform a particular persistent reservation service action.

[2] The parameter list shall be 24 bytes in length and the Parameter List Length field shall contain 24 (18h). If the Parameter List Length is not 24, the device server shall return a Check Condition status. The sense key shall be set to Illegal Request and the additional sense data shall be set to Parameter List Length Error.

[3] When processing the Persistent Reserve Out service actions, the device server shall increment the generation value as specified in 5.15.1.

The Persistent Reserve Out command Service Action codes are briefly described in Table 115.

Table 115: Persistent Reserve Out Service Action codes

Code	Name	Description
00h	Register	The Persistent Reserve Out command executing a Register service action registers a reservation key with a device server. See T10/1157-D, Section 5.3, for details.
01h	Reserve	The Persistent Reserve Out command performing a Reserve service action creates a persistent reservation having a specified scope and type. The scope and type of a persistent reservation are defined in notes [3] and [4] of Table 111. See T10/1157-D, Section 5.3, for details.
02h	Release	The Persistent Reserve Out command performing a Release service action removes an active persistent reservation held by the same initiator. See T10/1157-D, Section 5.3, for details.
03h	Clear	The Persistent Reserve Out command that successfully performs a Clear service action shall remove all persistent reservations for all initiators. See T10/1157-D, Section 5.3, for details.
04h	Preempt	Preempts persistent reservations from another initiator. See T10/1157-D, Section 5.3, for details.
05h	Preempt and Abort	Preempts persistent reservations from another initiator and aborts the task set for the preempted initiator. See T10/1157-D, Section 5.3, for details.
06h-1Fh	Reserved	

[4] See “Control byte,” paragraph 4.2.6.

5.16.1 Persistent Reserve Out parameter list

The parameter list required to perform the Persistent Reserve Out command are defined in Table 116. All fields shall be sent on all Persistent Reserve Out commands, even if the field is not required for the specified Service action and Scope values.

Table 116: Persistent Reserve Out parameter list

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

Notes.

[1] The Reservation Key field contains an 8-byte token provided by the application client to the device server to identify the initiator that is the source of the Persistent Reserve Out command. The device server shall verify that the Reservation key field in a Persistent Reserve Out command matches the registered reservation key for the initiator from which the command was received. If a Persistent Reserve Out command specifies a Reservation key field other than the reservation key registered for the initiator, the device server shall return a Reservation Conflict status. The reservation key of the initiator shall be valid for all Service action and Scope values. The obsolete field in Bytes 22 and 23 was used for extent reservations.

- [2] The Service Action Reservation Key field contains information needed for three service actions; the Register, Preempt, and Preempt and Clear service actions. For the Register service action, the Service Action Reservation key field contains the new reservation key to be registered. For the Preempt and Preempt and Clear service actions, the Service Action Reservation key field contains the reservation key of the persistent reservations that are being preempted. For the Preempt and Preempt and Clear service actions, failure of the Service Action Reservation key to match any registered reservation keys shall result in the device server returning a Reservation Conflict status. The Service Action Reservation key is ignored for all service actions except those described in this paragraph.
- [3] If the Scope is an Element reservation, the Scope-specific Address field shall contain the Element address, zero filled in the most significant bytes to fit the field. If the Service action is Register or Clear or if the Scope is a Logical Unit reservation, the Scope-specific address field shall be set to zero.
- [4] The Activate Persist Through Power Loss (APTPL) bit shall be valid only for the Register service action. In all other cases, the APTPL shall be ignored. Support for an APTPL bit equal to one is optional. If a device server that does not support the APTPL bit value of one receives that value in a Register service action, the device server shall return a Check Condition status. The sense key shall be set to Illegal Request and additional sense data shall be set to Invalid Field in Parameter List.

If the last valid APTPL bit value received by the device server is zero, the loss of power in the target shall release all persistent reservations and remove all reservation keys (see Table 115, register description). If the last valid APTPL bit value received by the device server is one, the logical unit shall retain all persistent reservations and all reservation keys for all initiators even if power is lost and later returned. The most recently received valid APTPL value from any initiator shall govern logical unit's behavior in the event of power loss.

Table 117 summarizes which fields are set by the application client and interpreted by the device server for each Service Action and Scope value. Two Persistent Reserve Out parameters are not summarized in Table 117; Reservation key and APTPL, since they are specified in Table 116.

Table 117: Persistent Reserve Out Service Actions and valid parameters

Service Action	Allowed Scope	Parameters		
		Type	Service Action Reservation Key	Element or Element Parameters
Register	ignored	ignored	valid	ignored
Reserve	LU	valid	ignored	ignored
Reserve	Extent	valid	ignored	Extent valid
Reserve	Element	valid	ignored	Element valid
Release	LU	valid	ignored	ignored
Release	Extent	valid	ignored	Extent valid
Release	Element	valid	ignored	Element valid
Clear	ignored	ignored	ignored	ignored
Preempt	LU	valid	valid	ignored
Preempt	Extent	valid	valid	Extent valid
Preempt	Element	valid	valid	Element valid
Preempt and clear	LU	valid	valid	ignored
Preempt and clear	Extent	valid	valid	Extent valid
Preempt and clear	Element	valid	valid	Element valid

5.17 Prefetch command (34h)

The Prefetch command (Table 118) 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 118: Prefetch command (34h)

Bit Byte	7	6	5	4	3	2	1	0	
0	Operation Code (34h)								
1	Logical Unit Number			Reserved			Immed [1]	RelAdr [2]	
2	Logical Block Address [2]								
3									(MSB)
4									
5									(LSB)
6	Reserved								
7	Transfer Length [3]								
8									(MSB)
9	Control [4]								
							(LSB)		

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] 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.
- [3] 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 paragraph 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 paragraph 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 paragraph 4.2.6) is one, the drive shall return Intermediate status.
- [4] See "Control byte," paragraph 4.2.6.

5.18 Prevent/Allow Medium Removal command (1Eh)

Not implemented by the drive. If received, the drive terminates with Check Condition status and sets an Illegal Request sense key.

5.19 Read command (08h)

The Read command (Table 119) requests that the drive transfer data to the initiator.

Table 119: Read command (08h)

Bit Byte	7	6	5	4	3	2	1	0
0	0	0	0	0	1	0	0	0
1	0	0	0	Logical Block Address (MSB) [2]				
	Logical Unit Number [1]							
2	Logical Block Address							
3	Logical Block Address (LSB)							
4	Transfer Length [3]							
5	0	0	0	0	0	0	Flag [4]	Link [4]

Notes.

- [1] The Logical Unit Number 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 5.10). 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 drive's Product Manual, Volume 1, 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 5.35) 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 below)
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.

5.20 Read Buffer command (3Ch)

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 120.

Table 120: Read Buffer command (3Ch)

Bit Byte	7	6	5	4	3	2	1	0	
0	0	0	1	1	1	1	0	0	
1	Logical Unit Number [1]			0	Reserved				Mode [2] [6]
2	Buffer ID [5]								
3	(MSB)								
4	Buffer Offset [5]								
5									(LSB)
6	(MSB)								
7	Allocation Length [3]								
8									(LSB)
9	0	0	0	0	0	0	Flag [4]	Link [4]	

Notes.

- [1] The Logical Unit Number must be zero.
- [2] Read Buffer Mode field. The interpretation of data bits 0-3 is given in the following table and referenced sections.

DB3	DB2	DB1	DB0	Mode definition
0	0	0	0	Read combined descriptor header and data (Section 5.20.1)
0	0	1	0	Read data (Section 5.20.2)
0	0	1	1	Read descriptor (Section 5.20.3)
0	1	0	0	Reserved
0	1	0	1	Reserved
0	1	1	0	Reserved
0	1	1	1	Reserved
1	0	1	0	Read data from echo buffer (Section 5.20.4) [6]
1	0	1	1	Echo buffer descriptor mode (Section 5.20.5) [6]
1	0	0	0	Reserved
1	0	0	1	Reserved
1	1	0	0	Reserved
:	:	:	:	Reserved
1	1	1	1	Reserved

- [3] Drive specific size also depends on whether cache is enabled as shown in Table 95 (RCD = 0) or not (RCD = 1). See individual drive's Product Manual, Volume 1, under list of SCSI commands supported.
- [4] See "Control byte," paragraph 4.2.6.
- [5] Not implemented by drives supported by this manual. Must be zero.
- [6] The Read data from echo buffer mode and Echo buffer descriptor mode do not apply to drives marketed before the release of Rev. J of this manual. Drives marketed after Rev. J may or may not support these two modes. See individual drive's Product Manual, Volume 1, for declaration of support.

5.20.1 Read Combined Descriptor Header and Data mode (0000b)

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 (see Table 39).

Table 121: Read Buffer header (0000b)

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

[1] The Buffer Capacity field specifies the total number of data bytes that are available in the drive's data buffer (see Sections 5.51.1 and 5.51.2). 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.

5.20.2 Read Data mode (0010b)

In this mode, the Data In phase contains buffer data only with no header. The Buffer ID and Buffer Offset fields are not used.

5.20.3 Read Buffer Descriptor mode (0011b)

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 Table 122.

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.

Table 122: Read Buffer Descriptor mode (0011b)

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

[1] 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 123.

Table 123: 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

- [2] The Buffer Capacity field specifies the total number of data bytes that are available in the drive's data buffer (see 5.51.1 and 5.51.2). 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.

5.20.4 Read Data from Echo Buffer mode (1010b)

In this mode, the target transfers data to the initiator from the echo buffer. The echo buffer shall transfer the same data as when the Write Buffer command with the mode field set to echo buffer was issued. The Buffer ID and Buffer Offset fields are ignored in this mode.

Note. A Write Buffer command with the mode field set to echo buffer should be sent prior to the Read Buffer command; otherwise the Read Buffer command may terminate with Check Condition: Illegal Request or may return indeterminate data.

The Read Buffer command shall return the same number of bytes of data as received in the prior Write Buffer command.

5.20.5 Echo Buffer Descriptor mode (1011b)

In this mode, a maximum of four bytes of Read Buffer descriptor information is returned. The target shall return the descriptor information for the echo buffer. If there is no echo buffer implemented, 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 Table 124.

Table 124: Echo Buffer Descriptor mode (1011b)

Bit Byte	7	6	5	4	3	2	1	0
0	Reserved							
1	(MSB)							
2	Buffer Capacity [1]							
3								

Notes.

- [1] The Buffer Capacity field shall return the size of the echo buffer in bytes aligned to a four-byte boundary.

5.21 Read Capacity command (25h)

The Read Capacity command (Table 125) provides a means for the initiator to request the capacity of the drive information.

Table 125: Read Capacity command (25h)

Bit Byte	7	6	5	4	3	2	1	0
0	0	0	1	0	0	1	0	1
1	Logical Unit Number [1]			0	0	0	0	RelAdr [2]
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	PMI [3]
9	0	0	0	0	0	0	Flag [4]	Link [4]

Notes.

- [1] The Logical Unit Number must be zero.
- [2] 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.
- [3] 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 of the logical unit. (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.) 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.
- [4] See "Control byte," paragraph 4.2.6.

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 126.

Table 126: 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)

5.22 Read Defect Data (10) command (37h)

The Read Defect Data (10) 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 127.

Table 127: Read Defect Data (10) command (37h)

Bit Byte	7	6	5	4	3	2	1	0
0	0	0	1	1	0	1	1	1
1	Logical Unit Number [1]			0	0	0	0	0
2	0	0	0	P	G	Defect List Format [3]		
				[2]		1	0	0 or 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 [5]	Link [5]

Notes.

[1] The Logical Unit Number 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	Return Defect List header only. Target returns only the Defect List header.
0	1	Return the growth "G" list only. This list reflects the growth or "G" list as defined in Section 5.5.
1	0	Return the manufacturer's original ETF list only. This list reflects the manufacturer's 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).
1	1	Return all lists. The returned list contains all of the requested drive's defect lists (i.e., P, G, C, and D) regardless of whether these lists have been reallocated by the drive.

[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 Unit command (indicated in Table 38).

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 creates 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 5.5). 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 128.

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 127).

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 128: Defect List Header Description

Bit Byte	7	6	5	4	3	2	1	0
0	0	0	0	0	0	0	0	0
	Reserved							
1	0	0	0	P	G	1	0	0 or 1
	Reserved			[1]		Defect List Format [2]		
2	Defect List Length (MSB) (Value is 8 times the number of defects) [3]							
3	Defect List Length (LSB) [3]							
4-n	Defect Descriptor Bytes [3]							

Note.

[1] The drive interprets the P and G bits (bits 4 and 3 of byte 1 of the CDB) as follows:

Bit P Bit G

0	0	Return Defect List header only. Target returns only the Defect List header.
0	1	Return the growth “G” list only. This list reflects the growth or “G” list as defined in Section 5.5.
1	0	Return the manufacturer’s original ETF list only. This list reflects the manufacturer’s 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).
1	1	Return all lists. The returned list contains all of the requested drive’s defect lists (i.e., P, G, C, and D) regardless of whether these lists have been reallocated by the drive.

[2] 1 0 0 defines a list in the Bytes from Index format.

1 0 1 defines a list in the Physical Sector format.

[3] If the P and G bits are 0 0, the Defect List Length will be zero and no Defect Descriptor Bytes are sent to the initiator. See Tables 46 and 47 for format of Defect Descriptor Bytes.

5.23 Read Defect Data (12) command (B7h)

The Read Defect Data (12) command (see Table 129) requests that the device server transfer the medium defect data to the application client.

If the logical unit is reserved, a reservation conflict shall occur when a Read Defect Data (12) command is received from an initiator other than the one holding a logical unit reservation. The command shall be rejected with Reservation Conflict status if the reservation conflict is due to a logical unit reservation. Read Defect Data (12) commands with a reservation conflict shall be terminated with Reservation Conflict status. The Read Defect Data (12) command shall not be evaluated for extent reservation conflicts (e.g., extent reservations do not conflict with the Read Defect Data (12) command).

Table 129: Read Defect Data (12) command (B7h)

Bit Byte	7	6	5	4	3	2	1	0
0	Operation Code (B7h)							
1	Logical Unit Number [1]			P List [2]	G List [2]	Defect List Format [3]		
2	Reserved							
3	Reserved							
4	Reserved							
5	Reserved							
6	(MSB)							
7	Allocation Length [4]							
8								
9								
10	Reserved							
11	Control [5]							

Notes.

[1] The Logical Unit Number 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 Return Defect List header only. Target returns only the Defect List header.
- 0 1 Return the growth “G” list only. This list reflects the growth or “G” list as defined in Section 5.5.
- 1 0 Return the manufacturer’s original ETF list only. This list reflects the manufacturer’s 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).
- 1 1 Return all lists. The returned list contains all of the requested drive’s defect lists (i.e., P, G, C, and D) regardless of whether these lists have been reallocated by the drive.

[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 Unit command (indicated in Table 38).

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 creates 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.

The Read Defect Data (12) list header (see Table 130) contains an eight byte header, followed by zero or more defect descriptors.

Table 130: Read Defect Data (12) list header

Bit Byte	7	6	5	4	3	2	1	0
0	Reserved							
1	Reserved			P List [1]	G List [1]	Defect List Format [2]		
2	Reserved							
3	Reserved							
4	(MSB)							
5	Defect List Length [3]							
6								
7								
Defect Descriptors [3]								
0								
n								

Notes.

[1] The drive interprets the P and G bits (bits 4 and 3 of byte 1 of the CDB) as follows:

Bit P	Bit G	
-------	-------	--

0	0	Return Defect List header only. Target returns only the Defect List header.
0	1	Return the growth “G” list only. This list reflects the growth or “G” list as defined in Section 5.5.
1	0	Return the manufacturer’s original ETF list only. This list reflects the manufacturer’s 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).
1	1	Return all lists. The returned list contains all of the requested drive’s defect lists (i.e., P, G, C, and D) regardless of whether these lists have been reallocated by the drive.

[2] 1 0 0 defines a list in the Bytes from Index format.

1 0 1 defines a list in the Physical Sector format.

[3] If the P and G bits are 0 0, the Defect List Length will be zero and no Defect Descriptor Bytes are sent to the initiator. See Tables 46 and 47 for format of Defect Descriptor Bytes.

5.24 Read Element Status command (B4h)

Not implemented. If received, a Check Condition status is sent.

5.25 Read Extended command (28h)

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 131.

This command operates the same as the Read command (see Section 5.19) 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 5.35) exists, and no data shall be read.

Table 131: Read Extended command (28h)

Bit Byte	7	6	5	4	3	2	1	0
0	0	0	1	0	1	0	0	0
1	Logical Unit Number [1]			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]

Notes.

- [1] The Logical Unit Number 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.
- 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 note [4]).
- [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.

5.26 Read Long command (3Eh)

The Read Long command (Table 132) 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.

Table 132: Read Long command (3Eh)

Bit Byte	7	6	5	4	3	2	1	0
0	0	0	1	1	1	1	1	0
1	Logical Unit Number [1]			0	0	0	Correct [2]	RelAdr [3]
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 [6]	Link [6]

Notes.

- [1] The Logical Unit Number 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 (01h). See individual drive's Product Manual, Volume 1, for a detailed discussion of the 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.

5.27 Reassign Blocks command (07h)

The Reassign Blocks command (Table 133) 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 5.10).

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

Table 133: Reassign Blocks command (07h)

Bit Byte	7	6	5	4	3	2	1	0
0	0	0	0	0	0	1	1	1
1	0	0	0	0	0	0	0	0
	Logical Unit Number [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	0
5	0	0	0	0	0	0	Flag [2]	Link [2]

Notes.

- [1] The Logical Unit Number must be zero.
- [2] See "Control byte," paragraph 4.2.6.

5.27.1 Reassign Blocks defect list

The Reassign Blocks defect list (Table 134) contains a four byte header followed by one or more Defect Descriptors. The length of each Defect Descriptor is four bytes.

Table 134: Reassign Blocks defect list

Bit Byte	7	6	5	4	3	2	1	0
Defect List Header								
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.

5.28 Rebuild command (81h)

The Rebuild command (Table 135) 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. For additional information about Xor commands, see Section 5.55.1.

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.

Table 135: Rebuild command (81h)

Bit Byte	7	6	5	4	3	2	1	0
0	Operation Code (81h)							
1	Reserved			DPO [1]	FUA [2]	IntData [3]	Port Control [4]	
2	(MSB) _____ Logical Block Address [5] _____ (LSB)							
3								
4								
5								
6	(MSB) _____ Rebuild Length [6] _____ (LSB)							
7								
8								
9								
10	(MSB) _____ Parameter List Length [7] _____ (LSB)							
11								
12								
13								
14	Reserved							
15	Control [8]							

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 136. 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 136: 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 target's choice.
11	The target transfers the data using one or more ports of the target's 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 137).
- [8] See "Control byte," paragraph 4.2.6.

5.28.1 Rebuild and Regenerate parameter data

The Rebuild and Regenerate parameter data is described in Table 137.

Table 137: Rebuild and Regenerate parameter data

Bit Byte	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 : 19	Source Descriptor (first) [3]							(LSB)
	: [3]							
16x – 12 16x + 3	Source Descriptor (last) [3]							
16x + 4 16x+y+3	Pad (if any) (y=length) [4]							
16x+y+z+4	(MSB)	Intermediate Data (if any) (z=length) [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 138 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 138: Source Descriptor format

Bit Byte	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.

5.29 Receive Diagnostic Results command (1Ch)

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 164), the data returned to the initiator is in the format shown in Table 140.

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 141 or 142, respectively.

All FRU and error code definitions are unique to this product and intended for factory/field maintenance personnel.

Table 139: Receive Diagnostic Results command (1Ch)

Bit Byte	7	6	5	4	3	2	1	0
0	0	0	0	1	1	1	0	0
1	Logical Unit Number [1]			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 [3]	Link [3]

Notes.

- [1] The Logical Unit Number 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 140: 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] A Field Replaceable Unit (FRU) 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] The Error Code is a two byte value that 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].
- [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 Mismatch While Doing Drive Diagnostics

- [5] Additional Vendor Unique codes (not available).

5.29.1 Supported Diagnostic Pages

If the Send Diagnostics command requests the Supported Diagnostics Pages list (PF bit = 1), the drive returns data in the format shown in Table 141 after receiving the Receive Diagnostic Results command. It lists all of the diagnostics pages supported by the drive.

Table 141: Supported Diagnostic Pages

Bit Byte	7	6	5	4	3	2	1	0
0	Page Code (00h)							
1	Reserved							
2	(MSB)							
3	Page Length (n-3) [1] (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 141) and 40h (Table 142).

5.29.2 Translate Address page (40h)

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 142).

Table 142: Translate Address page—Receive Diagnostic (40h)

Bit Byte	7	6	5	4	3	2	1	0
0	Page Code (40h) [1]							
1	Reserved							
2	(MSB)							
3	Page Length (000A) or (0002) [2] (LSB)							
4	Reserved				Supplied Format [3]			
5	RAREA [4]	ALTSEC [5]	ALTTK [6]	Reserved	Translated Format [7]			
6 : 13	Translated Address (if available) [8]							

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 166).

- [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 143: Address Field Logical Block Address Format

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB)							
:	Logical Block Address							
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 144: Address Field Physical Sector Address Format

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB)							
1	Cylinder Number							
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 166). 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 143 and 144.

5.30 Regenerate command (82h)

The Regenerate command (see Table 145) 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. For additional information about Xor commands, see Section 5.55.1.

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.

Table 145: Regenerate command (82h)

Bit Byte	7	6	5	4	3	2	1	0	
0	Operation Code (82h)								
1	Reserved			DPO [1]	FUA [2]	IntData [3]	Port Control [4]		
2	Logical Block Address [5]								
3									(MSB)
4									
5									(LSB)
6	Regenerate Length [6]								
7									(MSB)
8									
9									(LSB)
10	Parameter List Length [7]								
11									(MSB)
12									
13									(LSB)
14	Reserved								
15	Control [8]								

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 146. 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 146: 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 target's choice.
11	The target transfers the data using one or more ports of the target's choice.

- [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 137. 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.

5.31 Release (6) command (17h)

The Release (6) command (Table 147) 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. For additional information about Reservations, see Section 5.35.5.

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 147.

Table 147: Release (6) command (17h)

Bit Byte	7	6	5	4	3	2	1	0
0	0	0	0	1	0	1	1	1
1	Logical Unit Number [1]			3rd Party [2]	3rd Party Device ID [2]			Extent [3]
2	0	0	0	0	0	0	0	0
	Reservation Identification [4]							
3	0	0	0	0	0	0	0	0
	Reserved							
4	0	0	0	0	0	0	0	0
	Reserved							
5	0	0	0	0	0	0	Flag [5]	Link [5]

Notes.

- [1] The Logical Unit Number 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] The Extent bit must be zero if not supported by the drive. See individual drive's Product Manual, Volume 1, section "SCSI Interface commands supported."
- [4] Must be zero if not supported by the drive. See individual drive's Product Manual, Volume 1, section "SCSI Interface commands supported." (See also paragraphs 5.35.1 and 5.35.2 in this manual).
- [5] See "Control byte," paragraph 4.2.6.

5.31.1 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.

5.31.2 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 individual drive's Product Manual, Volume 1, section "SCSI Interface commands supported" for support of this command.

5.31.3 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 5.35.3).

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.

5.32 Release (10) command (57h)

The Reserve and Release commands provide the basic mechanism for contention resolution in multiple-initiator systems. The Release (10) command (Table 148) 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 148: Release (10) command (57h)

Bit Byte	7	6	5	4	3	2	1	0
0	Operation Code (57h)							
1	0	0	0	3rd Party [2]	Reserved			0 Extent [3]
2	Reservation Identification [4]							
3	3rd 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 individual drive's Product Manual, Volume 1, SCSI Interface commands supported section (see paragraph 5.35.1 of this volume).
- [5] See "Control byte," paragraph 4.2.6.

For additional explanation about the Release command, see Section 5.31 and Section 5.35.5.

5.33 Report LUNs command (A0h)

The Report LUNs command (see Table 149) requests that the peripheral device logical unit numbers of known logical units in the target be sent to the application client. The Report LUNs command shall return information about only those logical units to which commands may be sent. When the HiSupport bit is one (see Table 49, Standard Inquiry data format), the device server shall support the Report LUNs command. A SCSI device that is capable of supporting a LUN address other than zero should support the Report LUNs command on logical unit zero.

The Report LUNs command shall not be affected by reservations or persistent reservations.

Table 149: Report LUNs command (A0h)

Bit Byte	7	6	5	4	3	2	1	0
0	Operation Code (A0h)							
1	Reserved							
2	Reserved							
3	Reserved							
4	Reserved							
5	Reserved							
6	(MSB)							
7	Allocation Length [1]							
8								
9								
10	Reserved							
11	Control [2]							

Notes.

[1] The Allocation Length shall be at least 16 bytes. If the Allocation Length is less than 16 bytes, the device server shall return Check Condition status. The sense key shall be set to Illegal Request and the additional sense data shall be set to Invalid Field in CDB.

The Allocation Length is not sufficient to contain the logical unit number values for all configured logical units, the device server shall report as many logical unit number values as will fit in the specified Allocation Length. This shall not be considered an error.

[2] See “Control byte,” paragraph 4.2.6.

The device server shall report the logical unit numbers of configured logical units using the format shown in Table 150.

Table 150: Report LUNs parameter list format

Bit Byte	7	6	5	4	3	2	1	0	
0	(MSB)	LUN List Length (n-7) [1]							(LSB)
3									(LSB)
4	(MSB)	Reserved							(LSB)
7									(LSB)
	LUN List								
8	(MSB)	First LUN							(LSB)
15									(LSB)
		.							
		.							
		.							
n-7	(MSB)	Last LUN							(LSB)
n									(LSB)

Notes.

- [1] The LUN List Length shall contain the length in bytes of the LUN list that is available to be transferred. The LUN List Length is the number of logical unit numbers reported multiplied by eight. If the allocation length in the command descriptor block is too small to transfer information about all configured logical units, the LUN List Length value shall not be adjusted to reflect the truncation.

5.34 Request Sense command (03h)

The Request Sense command (Table 151) requests that the drive transfer sense data to the initiator in the format shown in Table 152. 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:

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

Table 151: Request Sense command (03h)

Bit Byte	7	6	5	4	3	2	1	0
0	0	0	0	0	0	0	1	1
1	Logical Unit Number			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 [1]	Link [1]

Notes.

[1] See “Control byte,” 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.

5.34.1 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 152. Numbers in parentheses in the sense data are references to notes following the table. A 1 or an 0 means the data bit is always logic 1 or logic 0, respectively, when sent by a drive.

Table 152: Drive Extended Sense Data Summary

Bit Byte	7	6	5	4	3	2	1	0
0	Validity Bit [1]	1	1	1	0	0	0	0 or 1
	Error Code [2]							
1	0	0	0	0	0	0	0	0
	Segment Number [3]							
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 [12]						
16								
17								
18–n	Product Unique Sense Data [13]							

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; 71h means a deferred error. (See paragraph 5.34.3.1 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 156. the code given in byte 12 provides additional clarification of errors. (See note [10] for related information.)
- [8] If the validity bit is 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 logi-

cal block. For example, if the Sense Key is Medium Error, it is the Logical Block Address of the failure block.

- [9] Additional Sense 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 157. 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] The Sense Key Specific Valid (SKSV) bit 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 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 156.

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 153. These fields point to illegal parameters in command descriptor blocks and data parameters sent by the initiator in the Data Out phase.

Table 153: Field Pointer Bytes

Bit Byte	7	6	5	4	3	2	1	0
15	SKSV	C/D	Reserved		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 154. These fields identify the actual number of retries used in attempting to recover from the error condition.

Table 154: Actual Retry Count Bytes

Bit Byte	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 individual drive's Product Manual, Volume 1.

Note. This field relates to the retry count fields specified within the error recovery parameters page of the Mode Select command (see Table 94).

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 155. These fields are only defined for the Format Unit command with the Immed bit set to one.

Table 155: Format Progress Indication Bytes

Bit Byte	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 individual drive's Product Manual, Volume 1, SCSI bus conditions and miscellaneous features supported section.

[13] Bytes 18-n are not presently used.

5.34.2 Sense Keys

Table 156 lists the Sense Keys in the extended sense data format that are used by the drive.

Table 156: 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 non-recovered 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 non-recoverable 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 4.6.5 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	Mismcompare—Indicates that the source data did not match the data read from the medium.

5.34.3 Extended Sense, Additional Sense and Additional Sense Qualifier codes

Table 157 lists the Extended Sense, Additional Sense and Additional Sense Qualifier codes.

Note. Table 157 is for reference only, as not all drives listed on the cover of this manual support all of the codes listed.

Table 157: 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

Table 157: Error codes for bytes 12 and 13 of Sense Data (values are in hexadecimal) (continued)

Byte 12	Byte 13	Description
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—Recommand Reassignment
18	07	Recovered Data With ECC—Data Rewritten
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
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
29	00	Power On, Reset, Or Bus Device Reset Occurred
29	01	Power-On Reset Occurred
29	02	SCSI Bus Reset Occurred
29	03	Bus Device Reset Function Occurred
29	04	Device Internal Reset Occurred
29	05	Transceiver Mode Changed To Single-Ended
29	06	Transceiver Mode Changed To LVD
2A	01	Mode Parameters Changed
2A	02	Log Parameters Changed
2A	03	Reservations preempted

Table 157: Error codes for bytes 12 and 13 of Sense Data (values are in hexadecimal) (continued)

Byte 12	Byte 13	Description
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

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

5.34.3.1 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.

5.35 Reserve (6) 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 158 shows the command structure.

Table 158: Reserve (6) command (16h)

Bit Byte	7	6	5	4	3	2	1	0
0	0	0	0	1	0	1	1	0
1	Logical Unit Number [1]			3rd Party 0 or 1 [2]	3rd Party Device ID [2]			Extent [3]
2	Reservation Identification [4]							
3	Extent List Length (MSB) [3]							
4	Extent List Length (LSB) [3]							
5							Flag [5]	Link [5]

Notes.

- [1] The Logical Unit Number 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 individual drive's Product Manual, Volume 1, section "SCSI Interface commands supported."
- [4] Reserve Identification applies only if Extent is supported. This field must be zero if not supported. Check with individual drive's Product Manual, Volume 1, section "SCSI Interface commands supported" (also see paragraphs 5.35.1 and 5.35.2 in this volume).
- [5] See "Control byte," paragraph 4.2.6.

5.35.1 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), an Inquiry command (which shall be executed), or a Request Sense command (which shall be executed), the command is rejected with Reservation Conflict status.

5.35.2 Extent Reservation

The drive may not support Extent reservations, in which case this bit must always be zero. See individual drive's Product Manual, Volume 1, 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.

5.35.3 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 5.31.3). This feature is intended for use in multiple-initiator systems that use the Copy command. Check individual drive's Product Manual, Volume 1, 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.

5.35.4 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.

5.35.5 Reservations

Reservations may be used to allow a device server to execute commands from a selected set of initiators. The device server rejects any commands from initiators outside the selected set by uniquely identifying initiators using protocol specific mechanisms.

Application clients may add or remove initiators from the selected set using reservation commands. If the application clients do not cooperate in the reservation protocol, data may be unexpectedly modified and deadlock conditions may occur.

The general description of reservations involves two groups of considerations:

- a. the type of reservation established, and
- b. the method used to establish, rescind, and manage the reservation.

Table 159: Reservation commands

Operation			
Code	Command Name	Section	Page
16h	Reserve (6)	5.35	200
56h	Reserve (10)	5.36	202
17h	Release (6)	5.31	187
57h	Release (10)	5.32	188
5Eh	Persistent Reserve In	5.15	153
5Fh	Persistent Reserve Out	5.16	158

The subject of Reservations is large and detailed, and therefore not covered entirely in this manual. It is recommended that the reader refer to ANSI document T10/1157-D, sections 5.3, 7.11, 7.12, 7.16, 7.17, 7.21, and 7.22, for a greater understanding of and details pertaining to the reservation process as needed for design purposes.

5.36 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 160) 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 160: Reserve (10) command (56h)

Bit Byte	7	6	5	4	3	2	1	0
0	Operation Code (56h)							
1	0	0	0	3rd Party [2]	Reserved			Extent [3]
2	Reservation Identification [4]							
3	3rd Party Device ID [2]							
4	Reserved							
5	Reserved							
6	Reserved							
7	(MSB)							
8	Extent List Length [3]							
	(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 individual drive's Product Manual, Volume 1.
- [4] Must be zero if not supported by the drive. See individual drive's Product Manual, Volume 1.
- [5] See "Control byte," paragraph 4.2.6.

For additional information about this command, see Section 5.35 and Section 5.35.5.

5.37 Rezero Unit command (01h)

The Rezero Unit command (Table 161) 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 Unit 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.

Table 161: Rezero Unit command (01h)

Bit Byte	7	6	5	4	3	2	1	0
0	0	0	0	0	0	0	0	1
1	Logical Unit Number [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]

Notes.

[1] The Logical Unit Number must be zero.

[2] See "Control byte," paragraph 4.2.6.

5.38 Search Data Equal command (31h)

Not implemented. If received, a Check Condition status is sent.

5.39 Search Data High command (30h)

Not implemented. If received, a Check Condition status is sent.

5.40 Search Data Low command (32h)

Not implemented. If received, a Check Condition status is sent.

5.41 Seek command (0Bh)

The Seek command (Table 162) 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.

Table 162: Seek command (0Bh)

Bit Byte	7	6	5	4	3	2	1	0
0	0	0	0	0	1	0	1	1
1	0 0 0 Logical Unit Number [1]			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 [3]	Link [3]

Notes.

- [1] The Logical Unit Number must be zero.
- [2] The maximum Logical Block Address that may be specified for a Seek command is defined in Read Capacity Data, Section 5.21.
- [3] See "Control byte," paragraph 4.2.6.

5.42 Seek Extended command (2Bh)

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 163.

Table 163: Seek Extended command (2Bh)

Bit Byte	7	6	5	4	3	2	1	0
0	0	0	1	0	1	0	1	1
1	Logical Unit Number [1]			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 [3]	Link [3]

Notes.

- [1] The Logical Unit Number must be zero.
- [2] This command operates the same as the Seek command (Section 5.41) except that a four byte Logical Block Address is specified.
- [3] See "Control byte," paragraph 4.2.6.

5.43 Send Diagnostic command (1Dh)

The Send Diagnostic command requests that the drive perform diagnostic tests on itself or perform other optional operations. Table 164 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 Diagnostic Results command and a subsequent Data In phase that returns data to the initiator. Using this latter procedure of Send Diagnostic/Receive Diagnostic Results commands, the initiator can ask the drive to return a list of optional operations it supports (Table 165) and then request a supported additional operation. The drive supports only the Translate Address operation (Table 166).

Table 164: Send Diagnostic command (1Dh)

Bit Byte	7	6	5	4	3	2	1	0
0	0	0	0	1	1	1	0	1
1	Logical Unit Number			PF [1]	Reserved	Self Test [2]	DEVOFL [3]	UnitOFL [4]
2	0	0	0	0	0	0	0	0
3	0	0	0	0	0	0	0	0
	Parameter List Length (MSB) [5]							
4	0	0	0	0	(see explanation below)			
	Parameter List Length (LSB) [5]							
5	0	0	0	0	0	0	Flag [6]	Link [6]

Notes.

[1] A Page Format (PF) bit of one specifies that the data returned by the subsequent Receive Diagnostic Results command shall use the diagnostic page format described in the ANSI SCSI-2 specification. See Tables 165 and 166, which show formats of the two pages supported by the drive.

[2] A Self Test bit of one directs the drive to complete its default self test. If the self test is requested, the Parameter List Length shall be set to zero and no data shall be transferred. If the self test successfully passes, the command shall be terminated with a Good Condition status; otherwise, the command shall be terminated with a Check Condition status and, if Extended Sense is implemented, the sense key shall be set to Hardware Error.

The drive performs only its default self test. This includes seeking to a reserved nonuser-accessible cylinder and writing, reading, and verifying data for each data surface.

A Self Test bit of zero instructs the drive to perform the operation called out in the parameter list in the Data Out following the Send Diagnostics Command Data Block. For the drive, the pages that may be listed in the parameter list are given in note [5].

[3] The SCSI Device Off Line (DEVOFL) bit of 1 enables diagnostic operations that may adversely affect operations to other Logic Units on the same target. This bit is not interpreted by the drive.

[4] A logical Unit Off Line (Unit OFL) bit of zero disables write operations on user medium or operations that affect user visible medium positioning. The bit is not interpreted by the drive.

[5] Parameter List Length specifies the length in bytes of the parameter list that shall be transferred during the Data Out phase.

The Parameter list length will be either 00h, 04h, or 0Eh (14 Decimal). Zero means no additional parameter pages to be sent with this command (this is not to be considered an error), 04h is the length value sent when the initiator sends Page 00h (see Table 165), and 0Eh is the value sent when Page 40h is sent (see Table 166). No other pages are supported by the drive. If the initiator transfers more bytes than the drive supports, the drive rejects the command.

[6] See "Control byte," paragraph 4.2.6.

5.43.1 Supported Diagnostic page—Send Diagnostic

The Supported Diagnostic page (Table 165) instructs the drive to make available the list of all supported diagnostic pages to be returned by a subsequent Receive Diagnostic Results command. The definition of this page for the Send Diagnostic command includes only the first four bytes (Receive Diagnostic version given in Table 141). If the Page Length field is not zero, the drive shall terminate the Send Diagnostic command with a Check Condition status. The sense key shall be set to Illegal Request with an additional sense code of Invalid Field In Parameter List.

Table 165: Supported Diagnostic page

Bit Byte	7	6	5	4	3	2	1	0
0	Page Code (00h)							
1	Reserved							
2	(MSB)							
3	Page Length (must be zero)							
								(LSB)

5.43.2 Translate Address page—Send Diagnostic

The Translate Address page (Table 166) 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 a 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 returned after the Receive Diagnostic Results command (see Table 139).

Table 166: Translate Address page

Bit Byte	7	6	5	4	3	2	1	0
0	Page Code (40h)							
1	Reserved							
2	(MSB)							
3	Page Length (000Ah)							
								(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 143 and Table 144.

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

5.44 Set Limits command (33h)

Not implemented. If received, a Check Condition status is sent.

5.45 Start/Stop Unit command (1Bh)

The Start/Stop Unit command requests that the target enable the logical unit for further operations (start) or stop spindle rotation (stop).

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
- Request Sense
- Inquiry
- Reserve
- Release
- Motor Start
- Receive Diagnostic Results
- Write Data Buffer
- Read Data Buffer

The remaining commands (see individual drive’s Product Manual, Volume 1, SCSI interface commands supported section, for a 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.

Table 167: Start/Stop Unit command (1Bh)

Bit Byte	7	6	5	4	3	2	1	0
0	0	0	0	1	1	0	1	1
1	Logical Unit Number			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 [3]	Link [3]

Notes.

- [1] An Immediate (Immed) bit of zero indicates that status shall be returned after the operation is completed. If the Immed bit is one, status is returned as soon as the operation is initiated.
- [2] When the Start bit is 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 individual drive’s Product Manual, Volume 1.
- [3] See “Control byte,” paragraph 4.2.6.

5.46 Synchronize Cache command (35h)

The Synchronize Cache command (Table 168) 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.

Table 168: Synchronize Cache command (35h)

Bit Byte	7	6	5	4	3	2	1	0
0	0	0	1	1	0	1	0	1
1	0 0 0 Logical Unit Number [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]							

- [1] The Logical Unit Number 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] 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 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 "Control Byte," paragraph 4.2.6.

5.47 Test Unit Ready command (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 169 for proper format.

Table 169: Test Unit Ready command (00h)

Bit Byte	7	6	5	4	3	2	1	0
0	0	0	0	0	0	0	0	0
1	Logical Unit Number [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]

Notes.

[1] The Logical Unit Number must be zero.

[2] See "Control byte," 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.

5.48 Verify command (2Fh)

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 170. The drive disconnects while this command is being executed if the initiator supports disconnect/reconnect.

Table 170: Verify command (2Fh)

Bit Byte	7	6	5	4	3	2	1	0
0	0	0	1	0	1	1	1	1
1	0 0		0	0	0	0	BYT CHK [3]	RelAdr [4]
	Logical Unit Number [1]		DPO [2]					
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 [7]	Link [7]

Notes.

- [1] The Logical Unit Number 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, the 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] 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 verify operation shall begin, if RelAdr is zero (see note [4]).
- [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.

5.49 Write command (0Ah)

The Write command (Table 171) requests that the drive write, to the medium, the data transferred by the initiator.

Table 171: Write command (0Ah)

Bit Byte	7	6	5	4	3	2	1	0
0	0	0	0	0	1	0	1	0
1	0 0 0 Logical Unit Number [1]			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 [4]	Link [4]

Notes.

- [1] The Logical Unit Number 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 5.10).

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 5.19).

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 5.35) 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 below)
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.

5.50 Write and Verify command (2Eh)

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.

Table 172: Write and Verify command (2Eh)

Bit Byte	7	6	5	4	3	2	1	0
0	0	0	1	0	1	1	1	0
1	0 0		0	0	0	0	BYT CHK [3]	RelAdr [4]
	Logical Unit Number [1]			DPO [2]				
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]

Notes.

- [1] The Logical Unit Number 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, the 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] 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 note [4]).
- [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.

5.51 Write Buffer command (3Bh)

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 microcode.

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] of Table 173.

Table 173: Write Buffer command (3Bh)

Bit Byte	7	6	5	4	3	2	1	0
0	0	0	1	1	1	0	1	1
1	Logical Unit Number [1]			Reserved	Mode [2]			
2	Buffer ID [5]							
3	(MSB)							
4	Buffer Offset [5]							
5								
6	(MSB)							
7	Byte Transfer Length [3]							
8								
9	0	0	0	0	0	0	Flag [4]	Link [4]

Notes.

[1] The Logical Unit Number must be zero.

[2] The Mode field is defined in the following table and in the referenced sections.

Check with individual drive's Product Manual, Volume 1, to see which modes are supported by the drive in question.

DB3	DB2	DB1	DB0	Mode definition
0	0	0	0	Write combined header and data (Section 5.51.1)
0	0	0	1	Reserved
0	0	1	0	Write data (Section 5.51.2)
0	1	0	0	Download microcode (Section 5.51.3)
0	1	0	1	Download microcode and save (Section 5.51.4)
0	1	1	0	Download microcode with offsets (Section 5.51.5)
0	1	1	1	Download microcode with offsets and save (Section 5.51.6)
1	0	1	0	Write data to echo buffer (Section 5.51.7)
1	0	0	0	Reserved
1	0	0	1	Reserved
1	0	1	1	Reserved
:	:	:	:	Reserved
1	1	1	1	Reserved

[3] Use a Read Data Buffer command with mode bits set to 0011b 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] Implemented only by drives that support modes 0110b and 0111b.

5.51.1 Combined Header and Data mode (0000b)

In this mode, data to be written to the drive's 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 four byte Write Buffer header (which always precedes the data) plus the data to be written to the data buffer as follows:

Table 174: Write Buffer header

Bit Byte	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							

5.51.2 Write Data Only mode (0010b)

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.

5.51.3 Download Microcode mode (0100b)

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 (0100b) 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.

5.51.4 Download Microcode and Save mode (0101b)

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 4.6.5) 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 (0101b), the command bytes of Table 173 are interpreted as shown below:

5.51.4.1 Buffer ID and Buffer Offset fields (CDB bytes 2-5)

These bytes are all zero.

5.51.4.2 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.

5.51.5 Download Microcode with Offsets (0110b)

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 (0110b) 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 4.6.5) 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.

5.51.6 Download Microcode with Offsets and Save mode (0111b)

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 (0111b) 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 and 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 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 4.6.5) 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.

5.51.7 Write Data to Echo Buffer mode (1010b)

In this mode, the target transfers data from the initiator and stores it in an echo buffer. An echo buffer is assigned in the same manner by the target as it would for a write operation. Data shall be sent aligned on four-byte boundaries. The Buffer ID and Buffer Offset fields are ignored in this mode.

Upon successful completion of a Write Buffer command, the data shall be preserved in the echo buffer (see Section 5.20.4) unless there is an intervening command to any logical unit, in which case it may be changed.

5.52 Write Extended command (2Ah)

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 175. Refer also to individual drive’s Product Manual, Volume 1, Write Caching section, for information on Write Cache Control.

Table 175: Write Extended command (2Ah)

Bit Byte	7	6	5	4	3	2	1	0
0	0	0	1	0	1	0	1	0
1	Logical Unit Number [1]			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]

Notes.

- [1] The Logical Unit Number 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 note [4]).
- [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 5.49) 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 5.35) 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.

5.53 Write Long command (3Fh)

The Write Long command (Table 176) 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.

Table 176: Write Long command (3Fh)

Bit Byte	7	6	5	4	3	2	1	0
0	0	0	1	1	1	1	1	1
1	Logical Unit Number [1]						RelAdr [2]	
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	Byte Transfer Length (MSB) [3]							
8	Byte Transfer Length (LSB) [3]							
9	0	0	0	0	0	0	Flag [4]	Link [4]

Notes.

- [1] The Logical Unit Number must be zero.
- [2] 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 logical block at which the write operation shall occur.
- [3] 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.)
- [4] See "Control byte," paragraph 4.2.6.

5.54 Write Same command (41h)

The Write Same command (Table 177) requests that the target write the single block of data transferred by the initiator to the medium multiple times.

Table 177: Write Same command (41h)

Bit Byte	7	6	5	4	3	2	1	0
0	Operation Code (41h)							
1	Logical Unit Number [1]			Reserved		PBdata [2]	LBdata [3]	RelAdr [4]
2	(MSB)							
:	Logical Block Address [5]							
5								
6	Reserved							
7	(MSB)							
:	Number of Blocks [6]							
8								
9	Control [7]							

Notes.

- [1] The Logical Unit Number must be zero.
- [2] A Physical Block Data (PBdata) 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 (LBdata) 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] 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.
- [5] 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.
- [6] The 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.

5.55 XDRead command (52h)

The XDRead command (Table 178) requests that the target transfer to the initiator the Xor data generated by an XDWrite or Regenerate command. For additional information about Xor commands, see Section 5.55.1.

Table 178: XDRead command (52h)

Bit Byte	7	6	5	4	3	2	1	0	
0	Operation Code (52h)								
1	Reserved								
2	Logical Block Address [1]								
3									(MSB)
4									
5									(LSB)
6	Reserved								
7	Transfer Length [1]								
									(MSB)
8									(LSB)
9	Control [2]								

Notes.

- [1] 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.
- [2] See “Control byte,” paragraph 4.2.6.

5.55.1 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. Table 179 lists the Xor commands and the section number in which each is described. Mode page 10h “Xor control mode page” is used in conjunction with these commands and is listed in Section 5.12.11 (Table 100) with the other mode page descriptions.

Table 179: Xor commands defined in this section

Operation Code	Command Name	Section	Page
50h	XDWrite	5.56	225
51h	XPWrite	5.58	228
52h	XDRead	5.55	223
80h	XDWrite Extended	5.57	226
81h	Rebuild	5.28	178
82h	Regenerate	5.30	185

Not all Seagate devices support the Xor commands. Where used by a particular model Seagate device, the individual drive’s Product Manual, Volume 1, 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 opera-

tions 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.

5.56 XWrite command (50h)

The XWrite command (see Table 180) 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.

The resulting Xor data is retained in the target's buffer until it is retrieved by an XRead command with starting logical block address and transfer length fields that match the starting logical block address and transfer length of this command. For additional information about Xor commands, see Section 5.55.1.

Table 180: XWrite command (50h)

Bit Byte	7	6	5	4	3	2	1	0
0	Operation Code (50h)							
1	Reserved			DPO [1]	FUA [2]	Disable Write [3]	Reserved	
2	(MSB) _____ Logical Block Address [4] _____ (LSB)							
3								
4								
5								
6	Reserved							
7	(MSB) _____ Transfer Length [5] _____ (LSB)							
8								
9	Control [6]							

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 XWrite 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.

5.57 XDWrite Extended command (80h)

The XDWrite Extended command (see Table 181) 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. For additional information about Xor commands, see Section 5.55.1.

Table 181: XDWrite Extended command (80h)

Bit Byte	7	6	5	4	3	2	1	0								
0	Operation Code (80h)															
1	Table Address [1]	Reserved		DPO [2]	FUA [3]	Disable Write [4]	Port Control [5]									
2	Logical Block Address [6]															
3									(MSB)							
4																
5									(LSB)							
6	Secondary Logical Block Address [7]															
7									(MSB)							
8																
9									(LSB)							
10	Transfer Length [7]															
11									(MSB)							
12																
13									(LSB)							
14	Secondary Address [1]															
15	Control [8]															

Notes.

[1] A Table Address bit of zero indicates that the Secondary Address field contains the target identifier of the target to which the Xor data is transferred. The Logical Unit Number 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 lookup table of SAM-compliant target identifiers. This lookup 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 182. 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 182: 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 target's choice.
11	The target transfers the data using one or more ports of the target's choice.

- [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 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. The Secondary Logical Block Address field value 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.

- [8] 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.

5.58 XPWrite command (51h)

The XPWrite command (see Table 183) 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. For additional information about Xor commands, see Section 5.55.1.

Table 183: XPWrite command (51h)

Bit Byte	7	6	5	4	3	2	1	0	
0	Operation Code (51h)								
1	Reserved			DPO [1]	FUA [2]	Reserved			
2	Logical Block Address [3]								
3									(MSB)
4									
5									(LSB)
6	Reserved								
7	Transfer Length [4]								
8									(MSB)
8									(LSB)
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.

6.0 Seagate Technology support services

Online Services

Internet

For online information about Seagate products, visit www.seagate.com or e-mail your disc or tape questions to:

Presales Support:

Disc: http://www.seagate.com/support/email/email_presales.shtml or DiscPresales@Seagate.com

Tape: http://www.seagate.com/support/email/email_tape_presales.shtml or
Tape_Sales_Support@Seagate.com

Technical Support:

Disc: http://www.seagate.com/support/email/email_disc_support.shtml or DiscSupport@Seagate.com

Tape: http://www.seagate.com/support/email/email_tape_support.shtml or TapeSupport@Seagate.com

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). This service is available worldwide.

Automated Services

SeaFONE[®] (1-800-SEAGATE) is Seagate's toll-free number (1-800-732-4283) to access our automated self-help services. Using a touch-tone phone, you can find answers to service phone numbers, commonly asked questions, troubleshooting tips and specifications for disc drives and tape drives 24 hours daily. International callers can reach this service by dialing +1-405-936-1234.

SeaFAX[®] is Seagate's automated FAX delivery system. Using a touch-tone phone, you can obtain technical support information by return FAX 24 hours daily. This service is available worldwide.

Presales Support

Our Presales Support staff can help you determine which Seagate products are best suited for your specific application or computer system.

Seagate Express

You can purchase select tape products and tape accessories through Seagate Express 24 hours daily by calling 1-800-531-0968 or by faxing your order to: +1-972-481-4812.

Technical Support

If you need help 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 help, you can talk to a Seagate technical support specialist. Before calling, note your system configuration and drive model number (STxxxxx).

SeaTDD[™] (+1-405-936-1687) is a telecommunications device for the deaf (TDD). You can send questions or comments 24 hours daily and exchange messages with a technical support specialist from 8:00 A.M. to 12:15 P.M. and 1:30 P.M. to 6:00 P.M. (central time) Monday through Friday.

Customer Service (CSO)

Warranty Repair

Seagate offers worldwide customer support for Seagate drives. Seagate direct OEM, Distribution and System Integrator customers should contact their Seagate service center representative for warranty information. Other customers should contact their place of purchase.

Authorized Service Centers

If you live outside the U.S., you can contact an Authorized Service Center for service or repair.

USA/Canada/Latin America Support Services

Presales Support

Disc: 1-877-271-3285 or +1-405-936-1210 FAX: +1-405-936-1683
Tape: 1-800-626-6637 or +1-714-641-2500 FAX: +1-714-641-2410

Technical Support (SeaFONE)

1-800-SEAGATE or +1-405-936-1234 (for specific product phone number)
 FAX: *Disc:* +1-405-936-1685; *Tape:* +1-405-936-1683

SeaFAX 1-800-SEAGATE
SeaTDD +1-405-936-1687
SeaBOARD *Disc:* +1-405-936-1600; *Tape:* +1-405-936-1630

Warranty Repair

USA/Canada 1-800-468-3472 FAX: +1-405-949-6740
 Latin America +1-405-949-7758 FAX: +1-405-949-6738

Authorized Service Centers

Brazil
MA Informatica +55-21-516-6649 FAX: +55-21-516-5280

Canada
Memofix +1-905-660-4936 FAX: +1-905-660-4951
Adtech +1-905-812-8099 or 1-800-624-9857 FAX: +1-905-812-7807

European Support Services

For European customer support, dial the toll-free number for your specific country for presales support, technical support, SeaFAX and warranty repair.

If your country is not listed here, dial our European call center at +31-20-316-7222 from 8:30 A.M. to 5:00 P.M. (European central time) Monday through Friday. The European call center is located in Amsterdam, The Netherlands.

Call Center

Austria 0 800-20 12 90
 Belgium 0 800-74 876
 Denmark 80 88 12 66
 France 0 800-90 90 52
 Germany 0 800-182 6831
 Ireland 1 800-55 21 22
 Italy 1 677 90 695
 Netherlands 0 800-732 4283
 Norway 800-113 91
 Poland 00 800-311 12 38
 Spain 900-98 31 24
 Sweden 0 207 90 073
 Switzerland 0 800-83 84 11
 Turkey 00 800-31 92 91 40
 United Kingdom 0 800-783 5177

SeaBOARD

Germany +49-89-1409331

Fax Services—All European Countries

Presales/Technical Support/Warranty Repair 31-20-653-3513

Africa/Middle East Support Services

For presales, technical support, warranty repair and FAX services in Africa and the Middle East, dial our European call center at +31-20-316-7222 from 8:30 a.m. to 5:00 p.m. (European central time) Monday through Friday, or send a FAX to +31-20-653-3513. The European call center is located in Amsterdam, The Netherlands.

Asia/Pacific Support Services

Technical Support

Australia	+61-2-9725-3366	FAX: +61-2-9725-4052
Hong Kong	+852-2368 9918	FAX: +852-2368 7173
Japan	—	FAX: +81-3-5462-2979
Singapore	+65-488-7584	FAX: +65-488-7528
Taiwan	+886-2-2514-2237	FAX: +886-2-2715-2923

SeaFAX

Australia	+61-2-9756-5170
-----------	-----------------

Warranty Repair

Japan	+81-3-5462-2904	FAX: +81-3-5462-2979
Asia/Pacific and Australia	+65-485-3595	FAX: +65-488-7503

Appendix A SCSI Configured Automatically (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 off so the device is merely SCAM tolerant.

SCAM protocol is described in greater detail in the SCAM Product Manual, part number 77767519.

Index

Numerics

10 byte commands 76, 77
 12 byte commands 76, 77
 3rd Party. see Third Party
 6 byte commands 76, 77

A

abort command 213
 abort message 47, 162
 Abort Prefetch
 Caching Parameters page 134, 135
 abort tag message 162
 abort task 29, 47
 abort task message 33
 abort task set 29
 aborted command 3, 67, 73
 sense key 195
 aborted task 33, 41
 ABPF. see Abort Prefetch
 ACA. see Auto Contingent Allegiance
 ACK. see Acknowledge
 acknowledge
 SCSI bus signal 12
 ACKQ/REQQ
 Drive Standard Inquiry Data 88, 89
 AckReqQ. see ACKQ/REQQ
 acoustics 145
 activate persist through power loss bit 154
 Activate Persist Through Power Loss
 Persistent Reserve Out parameter list 159, 160
 Active Notch
 Notch page 139
 active notch 140
 active pointer 49
 actual length 153
 actual number of retries 193
 Actual Retry Count
 Sense Key Specific 194
 Actual Retry Count Bytes
 Sense Key Specific 193
 adapter
 host 11
 adaptive algorithm 127
 adaptive read look-ahead 135
 Additional Length
 Diagnostic Data Bytes 182
 Drive Standard Inquiry Data 88, 89
 Persistent Reserve In parameter data 154, 155
 additional parameter 195
 additional sense 195
 data 158, 189
 error code 85
 qualifier code 196
 Additional Sense Code
 Extended Sense Data Format 192, 193
 additional sense code 67, 68, 73, 83, 103, 105,
 106, 110, 112, 116, 122, 127, 144, 146,
 148, 169, 171, 175, 177, 196, 207, 210,
 211, 216, 217, 218
 Additional Sense Code Qualifier
 Extended Sense Data Format 193
 additional sense data 158
 Additional Sense Length
 Extended Sense Data Format 192, 193
 address
 task 7
 Address to Translate
 Translate Address page 207, 208
 address valid bit 162
 ADR16
 Drive Standard Inquiry Data 88, 90
 ADR32
 Drive Standard Inquiry Data 88, 90
 AEN 199
 AERC. see Asynchronous Event Reporting Capability
 Allocation Length
 Inquiry command 87
 Log Sense command 110, 111
 Mode Sense (10) command 149, 150
 Mode Sense (6) command 119, 120
 Persistent Reserve In command 153
 Read Buffer command 164
 Read Defect Data (10) command 169
 Read Defect Data (12) command 171, 172
 Receive Diagnostic Results command 181
 Report LUNs command 189
 Request Sense command 191
 allocation length 46, 150, 155, 165, 166, 170, 192
 Alternate Sector
 Translate Address page—Receive Diagnostic
 183, 184
 Alternate Sectors per Zone
 Format Parameters page 128, 129
 Alternate Track
 Translate Address page—Receive Diagnostic
 183, 184
 Alternate Tracks per Volume
 Format Parameters page 128, 129
 Alternate Tracks per Zone
 Format Parameters page 128, 129
 ALTSEC. see Alternate Sector
 ALTTK. see Alternate Track
 ANSI standards 2
 ANSI-Approved Version
 Drive Standard Inquiry Data 88, 89
 application client 2, 3, 148, 155, 160, 201
 application layer 6
 APTPL. see Activate Persist Through Power Loss
 arbitrating system 18, 25
 arbitration delay 14
 arbitration phase 17
 array controller 223, 224

- ARRE. see Automatic Read Reallocation Enabled
 - ASA-1 firmware 91
 - ASA-2 firmware 91
 - ASCII Family Number
 - Device Behavior page 101
 - ASCII space characters
 - Firmware Numbers page 96
 - assertion period 14
 - Association
 - Identification Descriptor 94, 95
 - asynchronous conditions 23
 - asynchronous event 7
 - notification 3, 137, 199
 - report 68
 - reporting 68, 69, 148
 - reporting sense data 69
 - Asynchronous Event Reporting Capability
 - Drive Standard Inquiry Data 88, 89
 - asynchronous information transfer 21
 - ATN signal 41
 - ATN. see Attention
 - attention
 - condition 23, 24
 - SCSI bus signal 12
 - attribute 66
 - auto contingent allegiance 48
 - auto contingent allegiance 3, 46, 48, 66, 68
 - (NACA=1) 69
 - active 66
 - active status 48
 - attribute 67
 - command 3
 - condition 47, 70
 - message 33
 - task 66
 - Automatic Read Reallocation Enabled
 - Error Recovery page 123, 124
 - Automatic Write Reallocation Enabled
 - Error Recovery page 123, 124
 - autosense 43, 69
 - autosense delivery 68
 - AWRE. see Automatic Write Reallocation Enabled
- B**
- Basic Queuing
 - Drive Standard Inquiry Data 88, 89
 - Behavior Code
 - Device Behavior page 101
 - Behavior Code Version Number
 - Device Behavior page 101
 - Bit G 169, 170, 171, 172
 - Bit P 169, 170, 171, 172
 - Bit Pointer
 - Sense Key Specific 193
 - Bit Pointer Valid
 - Sense Key Specific 193
 - block descriptor 151
 - Block Descriptor Length
 - Mode Select (10) parameter list 117
 - Mode Select parameter list 113
 - Mode Sense (6) command 119
 - Mode Sense Data 121
 - Mode Sense Parameter List 151
 - block format 82, 85
 - Block Length
 - Mode Select (10) parameter list 117, 118
 - Mode Select parameter list 113, 114
 - Mode Sense Data 121, 122
 - Mode Sense Parameter List 151, 152
 - block length 151
 - block size 185
 - blocked 3
 - blocking boundary 3
 - Board Serial Number
 - Unit Serial Number page 92
 - boundary field 165
 - BPV. see Bit Pointer Valid
 - BQue. see Basic Queuing
 - BSY. see Busy
 - buffer 218, 225
 - Buffer Capacity
 - Echo Buffer Descriptor mode 166
 - Read Buffer Descriptor mode 165, 166
 - Read Buffer header 165
 - buffer capacity 216, 217
 - Buffer Empty Ratio
 - Disconnect/Reconnect Control page 126, 127
 - buffer empty ratio 213
 - buffer field 166
 - Buffer Full Ratio
 - Disconnect/Reconnect Control page 126
 - buffer full ratio set 162
 - Buffer ID
 - Read Buffer command 164
 - Write Buffer command 215, 216
 - buffer ID 165, 216, 217, 218
 - code 217
 - field 218
 - Buffer Offset
 - Read Buffer command 164
 - Write Buffer command 215, 216
 - buffer offset 216, 217, 218
 - field 165, 166
 - buffer segment size 127
 - buffer testing applications 165
 - buffered operation 199
 - buffered write operation 199
 - burst data error 133
 - bus
 - data 12
 - SCSI 11
 - bus clear delay 14
 - bus device reset 29, 36, 38, 119, 149
 - bus message 47, 110, 162, 200, 201
 - other ports 40

- bus free
 - delay 14
 - phase 17, 71
 - state 49, 50
 - Bus Inactivity Limit
 - Disconnect/Reconnect Control page 126, 127
 - bus inactivity limit 127
 - bus phase 13
 - bus reset message 110
 - bus set delay 14
 - bus settle delay 14
 - bus signal
 - SCSI 12
 - bus timing 13
 - busy 41, 50, 72
 - SCSI bus signal 12
 - busy status 20, 47, 199
 - Busy Timeout Period
 - Control Mode page 137, 138
 - BYT CHK. see Byte Check
 - byte 3
 - byte by byte compare 212
 - Byte Check
 - Verify command 212
 - Write and Verify command 214
 - Byte Transfer Length
 - Read Long command 175
 - Write Buffer command 215, 217
 - Write Long command 221
 - byte transfer length 216
 - Bytes from Index
 - Defect List Header Description 170
 - Read Defect Data (12) list header 172
 - bytes from index format 82, 85
- C**
- C list 171
 - C type defect 81
 - C/D. see Command Data
 - C/D. see Control/Data
 - cable skew delay 14
 - cables 8
 - cache 157, 199, 213, 214
 - cache enable 178, 185, 219
 - cache memory 161, 210
 - cache replacement algorithm 135
 - Cache Segment Size
 - Caching Parameters page 134, 136
 - cache statistics page 108
 - Caching Analysis Permitted
 - Caching Parameters page 134, 135
 - caching mode page 213
 - caching page 145
 - caching parameters 115
 - page 120, 134, 150
 - page descriptor 122
 - call 3
 - CAP. see Caching Analysis Permitted
 - capacity programming 114, 118, 121, 151, 152
 - catastrophic failure 73
 - CDB. see Command Descriptor Block
 - certification defect 81
 - certification operation 194
 - Change Definition command 47, 73, 78
 - changeable parameter 113, 117, 149
 - changeable recovery time limit 133
 - changeable value 115, 126, 129, 134, 138, 149
 - changed operating definition 73
 - characteristic
 - logical 17
 - check condition 41, 43, 46, 66, 69
 - check condition status 31, 33, 45, 47, 66, 67, 68, 69, 72, 73, 79, 83, 85, 86, 87, 102, 104, 106, 110, 116, 122, 124, 127, 133, 142, 144, 148, 152, 158, 160, 161, 163, 169, 170, 171, 172, 174, 175, 177, 178, 179, 185, 189, 191, 193, 199, 200, 203, 206, 207, 208, 209, 210, 211, 212, 213, 214, 216, 217, 218, 220, 221, 223, 226, 227
 - check data 223, 224
 - Clear
 - Service Action code 159
 - clear ACA task management function 67
 - clear queue 69, 213
 - clear queue message 162
 - clear service action 154, 155, 159, 160
 - clear task set 30, 47
 - client 3, 6
 - client-server 3
 - model 3
 - CmdQue. see Command Queuing
 - CmpLst. see Complete List
 - Code Set
 - Identification Descriptor 94
 - code value
 - reserved 43
 - Combined Header and Data mode
 - Write Buffer command 216
 - command 3
 - 6, 10, 12 bytes 76, 77
 - aborted 3
 - alphabetical listing 75, 76
 - completed 3
 - ended 4
 - faulting 4
 - for all devices 75
 - for direct access devices 75
 - group 0, 1, 2 75
 - linked 5, 33
 - listed by operation code 75, 77
 - overlapped 67
 - SCSI, summary 9
 - SCSI-3 standards 2
 - single, example 49
 - third party 7

- unlinked 8
- untagged 33, 72
- command block 221
- command byte 216
- command code 44
- command complete message 49, 50, 78
- command completion 43
- Command Data
 - Sense Key Specific 193
- command descriptor block 3, 43, 44, 45, 67, 73, 84, 169, 195, 215
 - control byte 48
 - linked 5
- command length 76, 77
- command phase 22, 127, 199
- command processing 66
- command queue 47
- Command Queuing
 - Drive Standard Inquiry Data 88, 90
- command sequence 199
- command sequence error 216, 217
- command set 72
- command structure
 - packetized 43
- command terminated 195
- command terminated status 48, 199
- command transfer length 127
- command type
 - all devices 76, 77
 - direct access devices 76, 77
- common access method
 - SCSI-3 standards 2
- compare 212
- Compare command 79
- Compile Date Code
 - ETF Log Date Code page 98
- Complete List 82
 - Format Unit command 80
- completed command 3
- completed task 3
- completion of the current command 199
- completion status 162, 213, 227
- condition
 - asynchronous 23
 - attention 23, 24
 - auto contingent allegiance 48
 - reset 23, 24
- condition met status 47
- confirmed protocol service 4
- Connect Time Limit
 - Disconnect/Reconnect Control page 126, 127
- connect time limit 127
- connection 127
- connectors 8
- contention resolution 200
- contiguous logical blocks 212, 219
- contingent allegiance (NACA=0) 69
- contingent allegiance condition 41, 67, 138
- continue task 30
- Control
 - Change Definition command 78
 - Mode Select (10) command 116
 - Mode Sense (10) command 149
 - Persistent Reserve In command 153, 154
 - Persistent Reserve Out command 158
 - Prefetch command 161
 - Read Defect Data (12) command 171
 - Rebuild command 178
 - Regenerate command 185
 - Release (10) command 188
 - Report LUNs command 189
 - Reserve (10) command 202
 - Synchronize Cache command 210
 - Write Same command 222
 - XDRead command 223
 - XDWrite command 225
 - XDWrite Extended command 226
 - XPWrite command 228
- control 217, 218
- control byte 43, 66, 67
- control field 46
- control memory space 216
- control mode page 66, 71, 72, 105, 106, 115, 120, 137, 145, 150
 - descriptor 122
- Control. see also Flag and Link
- control/data
 - SCSI bus signal 12
- controller failure 195
- Copy and Verify command 79
- Copy command 79, 201
- Copyright Notice
 - Drive Standard Inquiry Data 88, 90
- Correct
 - Read Long command 175
- correct bit 175
- correction 133
- Correction Span
 - Error Recovery page 123, 124
- CRC 212
- critical data 199
- critical data transfer operation 203
- cumulative value
 - equal to threshold value 106
 - greater than threshold value 106
 - not equal threshold value 106
 - notify of update 106
- cumulative Xor operation 223
- current block size 221
- current command 26, 199
- current cumulative parameter 105
- current error 199
- current error information 199
- current mode value 112
- Current Operating Definition
 - Implemented Operating Definition page 92, 93

current operating definition 93
 current parameter 129
 current persistent reservation 154
 current pointer 26
 current reservation 201
 current sector 222
 current task 4
 current threshold parameter 105
 current value 149
 cyclic redundancy check 4
 cylinder 129
 cylinder boundary 50, 167
 cylinder number 140
 cylinder skew 129
 Cylinder Skew Factor
 Format Parameters page 128, 129

D

D list 171
 D type defect 81
 data
 failed 199
 subset 199
 data buffer 166, 216
 internal 213
 data bus 12
 data byte 221
 Data Bytes per Physical Sector
 Format Parameters page 128, 129
 data counter 105
 data defect 81
 data error correction 133
 data in phase 22, 43, 46, 120, 165, 166, 172,
 181, 191, 206, 207
 data out phase 22, 43, 46, 112, 183, 193, 206,
 217, 218
 data parity 12
 data phase 22, 38, 127
 data protect sense key 195
 data read from the medium 195
 data release delay 14
 Data Strobe Offset Count
 Error Recovery page 123, 124
 data transfer 212, 214, 223
 synchronous 22
 Data Transfer Disconnect Control
 Disconnect/Reconnect Control page 126, 127
 data transfer disconnect control field 31
 codes 127
 data transfer rate 127
 date code page 91, 98
 DBD. see Disable Block Descriptors
 DCR. see Disable Correction
 DCRT. see Disable Certification
 deadlock condition 201
 default initialization pattern 84
 Default Number of Cache Segments
 Device Behavior page 101
 Default Operating Definition
 Implemented Operating Definition page 92
 default operating definition 93
 default self test 206
 default value 149
 defect
 certification 81
 data 81, 172
 growth 81
 primary 81
 defect data descriptor 170
 defect data header 170
 Defect Descriptor 85
 Reassign Blocks defect list 177
 defect descriptor 86
 Defect Descriptor Bytes 85, 86
 Defect List Header Description 170
 defect descriptor bytes 81, 172
 Defect Descriptors
 Read Defect Data (12) list header 172
 defect list 131, 177
 error 83
 format field 86
 header 82, 83, 85
 header description 170
 not found 83, 169
 Defect List Format 82, 85
 Defect List Header Description 170
 Format Unit command 80
 Read Defect Data (10) command 169
 Read Defect Data (12) command 171
 Read Defect Data (12) list header 172
 Defect List Length
 Defect List header 83, 84
 Defect List Header Description 170
 Read Defect Data (12) list header 172
 Reassign Blocks defect list 177
 Defect Logical Block Address
 Reassign Blocks defect list 177
 defect management 129
 defect management zone 128, 129
 defect sector number of FFFFFFFFh 86
 defective logical block address 177
 deferred error 72, 192, 199
 deferred error report 199
 deferred error sense information 199
 Definition Parameter
 Change Definition command 78
 delay 14, 15
 Delayed Motor Start
 Jumper Settings page 99
 Demand Read Retention Priority
 Caching Parameters page 134, 135
 Density Code
 Mode Select (10) parameter list 117, 118
 Mode Select parameter list 113, 114
 Mode Sense Data 121

- Mode Sense Parameter List 151
- density code 151
- deskew delay 15
- destination device 4
- detected error 41
- device
 - destination 4
 - direct-access 92
 - SCAM 19
 - SCSI 6
 - source 7
- device behavior page 91, 101
- device failure 195
- device identification page 91, 93, 94
 - example 96
- Device Off Line
 - Send Diagnostic command 206
- device server 2, 4, 153, 155, 157, 158, 159, 160, 189
- device service
 - request 4
 - response 4
- DEVOFL. see Device Off Line
- DExcpt. see Disable Exception Control
- Diagnostic Data Bytes
 - Receive Diagnostic Results command 182
- diagnostic data list 182
- Diagnostic Fault Codes
 - Vendor Unique 182
- diagnostic function 164
- diagnostic mode 215
- diagnostic operation 182, 206
- diagnostic page format 206
- diagnostic test 206
- differential 4
- differential sense
 - SCSI bus signal 12
- DIFFSNS. see Differential Sense
- DImm. see Disconnect Immediate
- direct-access device 92
- Disable Block Descriptors
 - Mode Sense (10) command 149
 - Mode Sense (6) command 119
- Disable Certification
 - Defect List header 83
- Disable Correction
 - Error Recovery page 123, 124
 - Verify Error Recovery page 132, 133
- Disable Exception Control
 - Informational Exceptions Control page 146
- Disable Page Out
 - Caching Parameters page 136
 - Read Extended command 173
 - Rebuild command 178
 - Regenerate command 185
 - Verify command 212
 - Write and Verify command 214
 - Write Extended command 219
- XDWrite command 225
- XDWrite Extended command 226
- XPWrite command 228
- disable page out bit 185
- Disable Page Out-Force Unit Access
 - Mode Sense Data 121, 122
 - Mode Sense Parameter List 151, 152
- Disable Prefetch Transfer Length
 - Caching Parameters page 134, 135
- Disable Primary
 - Defect List header 83
- Disable Queuing
 - Control Mode page 137, 138
- Disable Read-Ahead
 - Caching Parameters page 134, 136
- disable retries 203
- Disable Save
 - Log Parameter 105
- disable save 103
- Disable Saving Parameters
 - Defect List header 83, 84
- Disable Transfer on Error
 - Error Recovery page 123, 124
 - Verify Error Recovery page 132, 133
- Disable Update
 - Log Parameter 105
- Disable Write
 - XDWrite command 225
 - XDWrite Extended command 226
- disc media 49
- disc servo-mechanism off-track error 199
- disc track 129
- DISC. see Discontinuity
- disconnect 4, 30, 127, 213, 214
 - capability 49
 - example 50
- Disconnect Immediate
 - Disconnect/Reconnect Control page 126, 127
- disconnect message 41, 47, 50
- disconnect privilege 31, 47, 71
 - bit 41, 127
- Disconnect Time Limit
 - Disconnect/Reconnect Control page 126, 127
- disconnect/reconnect 212
 - control page 20, 126
 - control page descriptor 122
 - control parameters 115, 120, 150
 - mode page 31
- disconnection delay 15
- Discontinuity
 - Caching Parameters page 134, 135
- DiscPriv. see Disconnect Privilege
- Distribution Diskette Serial Number
 - Drive Standard Inquiry Data 88, 90
- domain 4
- dormant
 - task state 4
- double transition 4

download 216
 download microcode 218
 download microcode and save 216
 Download Microcode and Save mode
 Write Buffer command 216
 Download Microcode mode
 Write Buffer command 216
 Download Microcode with Offsets
 Write Buffer command 217
 Download Microcode with Offsets and Save mode
 Write Buffer command 217
 download microcode with offsets and save mode
 217, 218
 downloadable code 217
 DPO. see Disable Page Out
 DPRY. see Disable Primary
 DQue. see Disable Queuing
 DRA. see Disable Read-Ahead
 drive counter 105
 drive failure 146
 Drive ID
 Jumper Settings page 99
 Drive ID table
 Jumper Settings page 100
 drive nominal rotation rate 131
 drive select 13
 Drive Serial Number
 Drive Standard Inquiry Data 88, 90
 drive specific parameter 169, 212
 drive specific size 164
 Drive Standard Inquiry Data 88
 Drive Step Rate
 Rigid Drive Geometry Parameters page 130
 Drive Type
 Format Parameters page 129
 drive user accessible capacity 129
 driver 4
 drives capacity capability 114
 DS. see Delayed Motor Start
 DS. see Disable Save
 DSP. see Disable Saving Parameters
 DTDC. see Data Transfer Disconnect Control
 DTE. see Disable Transfer on Error
 DU. see Disable Update
 dual port A-B 89
 dual port SCSI 40
 duplicate tag 67

E

ECC. see Error Correction Code
 echo buffer 164, 166, 218
 echo buffer descriptor mode 164, 166
 EECA. see Enable Extended Contingent Allegiance
 EER. see Enable Early Recovery
 electrical description 8
 Element
 Persistent Reservation 156, 160
 element address 156
 element reservation 156, 160
 EMDP. see Enable Modify Data Pointers
 Enable Early Recovery
 Error Recovery page 123, 124
 Verify Error Recovery page 132, 133
 Enable Extended Contingent Allegiance
 Control Mode page 137, 138
 Enable Modify Data Pointers
 Disconnect/Reconnect Control page 126, 127
 Enable Threshold Comparison
 Log Parameter 105, 106
 Enable Vital Product Data
 Inquiry command 87
 enable vital product data 91
 Enable Warning
 Informational Exceptions Control page 146
 enabled
 task state 4
 Enclosure Services
 Drive Standard Inquiry Data 88, 89
 EncServ. see Enclosure Services
 End of Medium
 Extended Sense Data Format 192
 end of medium partition 195
 ended command 4
 Ending Boundary
 Notch page 139, 140
 entire unit release 187
 EOM. see End of Medium
 equal comparison 195
 equal sense key 195
 error 66
 category, general 192
 current 199
 deferred 192, 199
 nonfatal 191
 non-recoverable 6, 195
 recovered 124, 162
 unrecoverable 124, 162
 unrecovered 199
 error and exception condition 43
 Error Code
 Diagnostic Data Bytes 182
 Extended Sense Data Format 192
 error code 199
 definition 181, 193
 sense data 196
 error condition 192, 193, 194, 199
 error corrected
 with possible delays code 107
 without substantial delay code 107
 error corrected code
 checksum 107
 error correction code 176, 212
 byte 175, 221
 correction 124, 175
 maximum 124, 133

- on-the-fly 107, 124, 133
- error counter page 107
- error event
 - recoverable 108
- error handling parameter 199
- error log 199
- error logging 137
- error recovery 124, 133, 175
 - flag 133
- error recovery page 123, 175
 - descriptor 122
- error recovery parameters 115, 120, 148, 150
- error recovery procedure
 - internal 213
- errors not related to any command 199
- establish reservation 201
- ETC. see Enable Threshold Comparison
- ETF list 81, 169, 170, 171
- ETF Log Date In ASCII
 - ETF Log Date Code page 98
- ETF Log Date in ASCII
 - ETF Log Date Code page 98
- EUI-64 identifier 96
- event
 - asynchronous 7
 - reset 6
- EVPD. see Enable Vital Product Data
- EWasc. see Enable Warning
- exception condition 4, 66, 194, 199
- Exclusive Access
 - Persistent Reservation 157
- execute command 213
- expected 8
- extended contingent allegiance 137
- extended message 33, 34
- extended sense 163, 174, 175, 206, 213, 216
 - codes 196
 - data byte 175
 - data format 191, 192, 195
 - incorrect length indicator bit 221
 - information byte 213
- Extent
 - Persistent Reservation 156, 160
 - Release (10) command 188
 - Release (6) command 187
 - Reserve (10) command 202
 - Reserve (6) command 200
- extent bit 200
- Extent Length
 - Persistent Reserve Out parameter list 159
- Extent List Length
 - Reserve (10) command 202
 - Reserve (6) command 200
- extent release 187
 - option 188
- extent reservation 159, 200
- external system intervention 199

F

- factory log page 109
- factory/field maintenance 181
- failed block 178
- failed data 199
- failed function 199
- failure prediction threshold exceeded 146
- fast assertion period 15
- fast cable skew delay 15
- fast deskew delay 15
- fast hold time 16
- fast negation period 16
- fast synchronous transfer option 15
- faulted initiator 4
- faulted task 66
- faulted task set 4
- faulting command 4
- faulting task 4
- FC-PH Name_Identifier 95
- fibre channel protocol 43
- field
 - maximum-value 72
 - minimum-value 72
 - reserved 43
 - unsupported 149
- Field Pointer
 - Sense Key Specific 193
- Field Pointer Bytes
 - Sense Key Specific 193
- Field Replaceable Unit
 - Diagnostic Data Bytes 182
- field replaceable unit 181
- Filemark
 - Extended Sense Data Format 192
- firmware error sense key 195
- firmware numbers page 91, 96
- First Reservation Key
 - Persistent Reserve In parameter data 154
- Flag
 - Format Unit command 80
 - Inquiry command 87
 - Log Select command 102
 - Log Sense command 110
 - Mode Select (6) command 112
 - Mode Sense (6) command 119
 - Read Buffer command 164
 - Read Capacity command 167
 - Read command 162
 - Read Defect Data (10) command 169
 - Read Extended command 173
 - Read Long command 175
 - Reassign Blocks command 176
 - Receive Diagnostic Results command 181
 - Release (6) command 187
 - Request Sense command 191
 - Reserve (6) command 200
 - Rezero Unit command 203

- Seek command 204
- Seek Extended command 205
- Send Diagnostic command 206
- Start/Stop Unit command 209
- Test Unit Ready command 211
- Verify command 212
- Write and Verify command 214
- Write Buffer command 215
- Write command 213
- Write Extended command 219
- Write Long command 221
- flag bit 46
- Flag. see also Control
- FmtData. see Format Data
- Force Sequential Write
 - Caching Parameters page 134, 136
- Force Unit Access
 - Caching Parameters page 136
 - Read Extended command 173
 - Rebuild command 178
 - Regenerate command 185
 - Write Extended command 219
 - XDWrite command 225
 - XDWrite Extended command 226
 - XPWrite command 228
- force unit access bit 185, 219, 225, 226, 228
- format 114
- format command 122, 152, 169, 172
- Format Data 82
 - Format Unit command 80
- format function 81, 152
- format operation 194
- Format Options Valid
 - Defect List header 83
- format parameters 115, 120, 150
 - page 128
 - page descriptor 122
- format progress 194
- Format Progress Indication Bytes
 - Sense Key Specific 194
- format unit
 - parameter definition 81, 82
 - parameter list 83, 84
- Format Unit command 80, 81, 124, 128, 129, 169, 170, 194, 195
- format verify phase 82
- FOV. see Format Options Valid
- FRU. see Field Replaceable Unit
- FSW. see Force Sequential Write
- FUA. see Force Unit Access
- function
 - failed 199
- function complete 4

G

- G bit 169, 170, 171, 172
- G list 171

- G type defect 81
- general error category 192
- general message protocol 27
- generate no sense 148
- generate unit attention 148
- Generation
 - Persistent Reserve In parameter data 154, 155
- Global Logging Target Save Disable
 - Control Mode page 137, 138
- glossary 3
- GLTSD. see Global Logging Target Save Disable
- good status 47, 112, 119, 149, 161, 178, 185, 187, 188, 199, 206, 211, 217, 218, 219, 225, 226
- group 0, 1, 2 commands 75
- group code field 44
- growth defect 81
- growth list 172

H

- hard reset 4, 90, 110, 119, 149
- hard reset condition 47, 200, 201
- Hard Sectoring
 - Format Parameters page 128, 129
- hardware error 177, 193, 206
- hardware error sense key 195
- head number 140
- head of queue 33
- head of queue task tag 32
- head of task queue tag message 71, 72
- Head Offset Count
 - Error Recovery page 123, 124
- hierarchial support bit 189
- Hierarchical Support
 - Drive Standard Inquiry Data 88, 89
- HiSupport. see Hierarchical Support
- hold time 15
- host adapter 11, 145
- HSEC. see Hard Sectoring

I

- I T L nexus 5
- I T L Q nexus 5
- I T nexus 5
- I/O operation 5
- I/O process 5, 17, 32, 39, 41, 138
 - termination, unsuccessful 17
- I/O system 6
- I/O. see also Input/Output
- IC. see Initiator Control
- ID bits
 - SCSI 11
- Identification Descriptor
 - Device Identification page 94
- identification descriptor 96
- Identification Descriptor List
 - Device Identification page 94

- Identifier
 - Identification Descriptor 94, 95
- Identifier Length
 - Identification Descriptor 94, 95
- Identifier Type
 - Identification Descriptor 94, 95
- identify 30, 31
- identify message 41, 43
- Idle
 - Power Condition page 141
- Idle Condition Timer
 - Power Condition page 141
- IEEE company_id.identifier field 95
- IEEE extended unique identifier 95
- ignore field definition 39
- ignore wide residue message 39
- IL bit 144
- ILI. see Incorrect Length Indicator
- illegal field 175
 - in CDB 221
 - in parameter list 127
- illegal parameter 193, 195
- illegal request 31, 43, 46, 79, 85, 103, 104, 106, 110, 116, 122, 127, 144, 158, 160, 161, 163, 166, 174, 175, 179, 185, 189, 193, 200, 207, 210, 216, 217, 218, 221, 223, 226
 - invalid command operation code 142
- illegal request sense key 72, 195
- ILQ nexus 32
- Immed. see Immediate
- Immediate
 - Defect List header 83, 84
 - Prefetch command 161
 - Start/Stop Unit command 209
 - Synchronize Cache command 210
- immediate bit 194, 199
- implementation 4
 - option 4
- implementation-specific 4
- implemented operating definition page 91, 92
- implied seek 212
- in transit 5
- inactive 141
- incorrect initiator connection 73
- Incorrect Length Indicator
 - Extended Sense Data Format 192
- incorrect length indicator
 - bit 175
- incorrect length indicator bit 195
- incorrect logical unit selection 67
- indeterminate data 166
- information
 - media 5
 - suspended 7
- Information Byte
 - Extended Sense Data Format 192
- information transfer
 - asynchronous 21
 - phase 13
- informational exception condition 148
- informational exceptions control page 115, 120, 146, 150
 - descriptor 122
- initialization operation 194
- Initialization Pattern
 - Defect List header 83, 84
 - Initialization Pattern Descriptor 84
- initialization pattern descriptor 84
- Initialization Pattern Length
 - Initialization Pattern Descriptor 84, 85
- Initialization Pattern Modifier
 - Initialization Pattern Descriptor 84
- Initialization Pattern Type
 - Initialization Pattern Descriptor 85
- initiator 2, 5, 11, 43, 49
 - causing 199
 - faulted 4
 - multiple 2, 11
 - single 11
 - temporary 178
- initiator connection
 - incorrect 73
- Initiator Control
 - Caching Parameters page 134, 135
- initiator detected error 31, 41
- initiator recovery action 48
- input/output
 - SCSI bus signal 12
- input/output. see also I/O
- inquiry 45, 46, 209
- Inquiry command 70, 72, 78, 87, 200
- inquiry data 46, 144
- Inquiry Length
 - Unit Attention Parameters page 144
- insufficient registration resources 158
- IntData. see Intermediate Data
- interconnect
 - physical 6
 - SCSI-3 standards 2
 - subsystem 5
- interface characteristics 8
- interface requirements 1
- Interleave
 - Format Parameters page 128, 129
 - Format Unit command 80
- interleave value 80
- Intermediate Data
 - Rebuild and Regenerate parameter data 179, 180
 - Rebuild command 178
 - Regenerate command 185
- intermediate status 48, 161
- intermediate-condition met status 48, 161
- internal data buffer 213
- internal error recovery procedure 213

- internal Xor operation 224
 - international standard 216
 - intersymbol interference (ISI) 5
 - Interval Timer
 - Informational Exceptions Control page 146, 147
 - invalid field in CDB 103, 110, 158, 175, 179, 185, 189, 210, 217, 218, 223, 226
 - invalid field in parameter list 85, 103, 104, 106, 116, 122, 144, 160, 207
 - invalid identify message 195
 - invalid logical block address 163, 174, 213, 220
 - invalid message error 41, 66
 - invalid parameter 43, 195
 - invalid task tag 33
 - IP Modifier. see Initialization Pattern Modifier
 - IP. see Initialization Pattern
 - IT nexus 7
 - ITL nexus 7, 32, 71
 - ITLQ nexus 7, 33, 71
 - ITX nexus 138
- J**
- jumper settings page 91, 99
- K**
- keyword 8
- L**
- Landing Zone Cylinder
 - Rigid Drive Geometry Parameters page 130
 - Last Reservation Key
 - Persistent Reserve In parameter data 154
 - layer 5
 - LBCSS. see Logical Block Cache Segment Size
 - LBdata. see Logical Block Data
 - LC 12
 - Link
 - Format Unit command 80
 - Inquiry command 87
 - Log Select command 102
 - Log Sense command 110
 - Mode Select (6) command 112
 - Mode Sense (6) command 119
 - Read Buffer command 164
 - Read Capacity command 167
 - Read command 162
 - Read Defect Data (10) command 169
 - Read Extended command 173
 - Read Long command 175
 - Reassign Blocks command 176
 - Receive Diagnostic Results command 181
 - Release (6) command 187
 - Request Sense command 191
 - Reserve (6) command 200
 - Rezero Unit command 203
 - Seek command 204
 - Seek Extended command 205
 - Send Diagnostic command 206
 - Start/Stop Unit command 209
 - Test Unit Ready command 211
 - Verify command 212
 - Write and Verify command 214
 - Write Buffer command 215
 - Write command 213
 - Write Extended command 219
 - Write Long command 221
 - Link. see also Control
 - linked CDB 5
 - Linked Command
 - Drive Standard Inquiry Data 88, 90
 - linked command 5, 33, 49, 71
 - complete 33, 46
 - List Parameter
 - Log Parameter 105, 106
 - Lock-unlock Cache command 101
 - Log Current Cumulative Values
 - Page Control field 103
 - Log Current Threshold Values
 - Page Control field 103
 - Log Default Cumulative Values
 - Page Control field 103
 - Log Default Threshold Values
 - Page Control field 103
 - Log Errors
 - Informational Exceptions Control page 146, 147
 - log page 102
 - codes 107
 - format 104
 - Log Parameter 105
 - Log Page format 104
 - log parameter page 102
 - Log Parameter Structure
 - Log Page format 104
 - log parameters changed 102
 - log save timer 203
 - Log Select command 102, 104
 - Log Sense command 104, 106, 110
 - LogErr. see Log Errors
 - logging information table 135
 - logical block 129, 214
 - Logical Block Address
 - Prefetch command 161
 - Read Capacity command 167
 - Read command 162
 - Read Extended command 173, 174
 - Read Long command 175
 - Rebuild command 178, 179
 - Regenerate command 185, 186
 - Seek command 204
 - Seek Extended command 205
 - Synchronize Cache command 210
 - Verify command 212
 - Write and Verify command 214
 - Write command 213

- Write Extended command 219
 - Write Long command 221
 - Write Same command 222
 - XDRead command 223
 - XDWrite command 225
 - XDWrite Extended command 226, 227
 - XPWrite command 228
 - logical block address 45, 84, 140, 173, 185, 192, 193, 203, 227
 - defective 177
 - format 184
 - Logical Block Cache Segment Size
 - Caching Parameters page 134, 136
 - logical block concept 45
 - Logical Block Data
 - Write Same command 222
 - logical block format 184
 - logical characteristics 17
 - Logical or Physical Notch
 - Notch page 139
 - Logical Unit
 - Persistent Reservation 156, 160
 - logical unit 5, 217, 218
 - not supported 31
 - option 5
 - Logical Unit Number
 - Change Definition command 78
 - Format Unit command 80
 - Inquiry command 87
 - Log Select command 102
 - Mode Select (10) command 116
 - Mode Select (6) command 112
 - Mode Sense (10) command 149
 - Mode Sense (6) command 119
 - Prefetch command 161
 - Read Buffer command 164
 - Read Capacity command 167
 - Read command 162
 - Read Defect Data (10) command 169
 - Read Defect Data (12) command 171
 - Read Extended command 173
 - Read Long command 175
 - Reassign Blocks command 176
 - Receive Diagnostic Results command 181
 - Release (10) command 188
 - Release (6) command 187
 - Request Sense command 191
 - Reserve (10) command 202
 - Reserve (6) command 200
 - Rezero Unit command 203
 - Seek command 204
 - Seek Extended command 205
 - Send Diagnostic command 206
 - Start/Stop Unit command 209
 - Synchronize Cache command 210
 - Test Unit Ready command 211
 - Verify command 212
 - Write and Verify command 214
 - Write Buffer command 215
 - Write command 213
 - Write Extended command 219
 - Write Long command 221
 - Write Same command 222
 - logical unit number 5, 27, 45
 - address 189
 - target 31
 - target routine number 31
 - Logical Unit Number List
 - Report LUNs parameter list format 190
 - Logical Unit Number List Length
 - Report LUNs parameter list format 190
 - logical unit persistent reservation 156
 - logical unit release 187
 - logical unit reservation 171, 200
 - logical unit reset 70
 - logical unit target 31
 - lookup table 226
 - low voltage differential 12
 - lower level protocol 5
 - lowest priority 214
 - low-level hardware parameter 72
 - LP. see List Parameter
 - LPN. see Logical or Physical Notch
 - LUN. see Logical Unit Number
 - LUNTAR. see Logical Unit Number Target
 - LUNTRN. see Logical Unit Number Target Routine Number
 - LVD. see Low Voltage Differential
 - LW 12
- ## M
- magnetic disc 92
 - maintenance 181
 - manage reservation 201
 - mandatory 5, 8
 - Maximum Burst Size
 - Disconnect/Reconnect Control page 126, 127
 - maximum burst size 127
 - Maximum Interleave
 - Device Behavior page 101
 - Maximum Number of Notches
 - Notch page 139
 - Maximum Prefetch
 - Caching Parameters page 134, 135
 - Maximum Prefetch Ceiling
 - Caching Parameters page 134, 135
 - Maximum Rebuild Read Size
 - Xor Control Mode page 142
 - Maximum Regenerate Size
 - Xor Control Mode page 142
 - Maximum Xor Write Size
 - Xor Control Mode page 142
 - maximum-value field 72
 - may 8
 - MChangr. see Medium Changer

- media
 - disc 49
- media information 5
- medium access command 211
- Medium Changer
 - Drive Standard Inquiry Data 88, 89
- medium defect data 169
- medium error 83, 124, 163, 169, 174, 177, 193
- medium error sense key 195
- Medium Rotation Rate
 - Rigid Drive Geometry Parameters page 131
- Medium Type
 - Mode Select (10) parameter list 117
 - Mode Select parameter list 113
 - Mode Sense Data 121
 - Mode Sense Parameter List 151
- medium verification 212
- memory
 - non-volatile 84, 87, 105
- message
 - code 28
 - detail 29
 - extended 34
 - one and two byte 29
 - SCSI bus signal 12
 - SCSI, summary 9, 10
- message exception conditions 40
- message in phase 23
- message out phase 23
- message parity error 32, 37
- message phase 23
- message reject 30, 31, 32, 35, 38, 41
- message reject message 138
- message system 27
- Method of Reporting Informational Exceptions
 - field format 148
 - Informational Exceptions Control page 146, 147
- MF. see Multiplication Factor
- microcode 217
- microcode has been changed 216, 217, 218
- Minimum Prefetch
 - Caching Parameters page 134, 135
- minimum-value field 72
- miscellaneous features supported 194
- miscompare 212, 214
- miscompare sense key 195
- Mode
 - Read Buffer command 164
 - Write Buffer command 215
- mode page 144, 223
- mode parameter 113, 120, 144, 151, 199
 - byte 149
- Mode Parameters
 - Mode Select page descriptor header 114
 - Mode Sense page descriptor header 122
- mode parameters changed 112
- Mode Select (10) command 116
- mode select (10) parameter list 116, 117
- Mode Select (6) command 112
- Mode Select command 69, 72, 115, 119, 126, 129, 131, 133, 134, 139, 140, 141, 143, 144, 145, 147, 149, 152, 194, 195, 199
- mode select condition 70
- mode select data 162
- Mode Select Page 8 178, 185, 219, 225, 226, 228
- mode select page descriptor 114
- Mode Select Page Headers
 - Mode Select (10) parameter list 117
 - Mode Select parameter list 113
- mode select parameter list 112, 113
- mode select parameters 69
- mode select savable parameters 84
- mode select/mode sense page 114
- mode sense 46
- Mode Sense (10) command 118, 149
- Mode Sense (6) command 119
- mode sense block descriptor 129
- Mode Sense command 45, 72, 112, 114, 116, 117, 129, 131, 133, 139, 144, 145, 152
- Mode Sense data 121, 126, 129, 133, 141, 143, 145, 147, 150
- Mode Sense Page 8 225, 226, 228
- mode sense page codes supported 115
- mode sense page descriptor header 122
- Mode Sense Page Headers
 - Mode Sense Data 121
- Mode Sense Page Headers and Their Parameters
 - Mode Sense Parameter List 151
- mode sense parameter list 151
- mode setting 195
- model
 - reference 6
- modify data pointer 26, 32
 - message 34, 127
- Motor Start
 - Jumper Settings page 99
- motor start 209
- Move Medium command 152
- MR head 83
- MRIE. see Method of Reporting Informational Exceptions
- MS. see Motor Start
- MSG. see Message
- Multi Port
 - Drive Standard Inquiry Data 88, 89
- multidrop 5
- multi-initiator environment 41
- multimode
 - SCSI bus signal 12
- multimode single-ended 5
- MultiP. see Multi Port
- multiple command buffer 199
- multiple deferred errors 199
- multiple I/O processes 32
- multiple initiator 11
- multiple initiator system 200

- multiple initiators 2
- multiple locks 210
- multiple persistent reservations 154
- multiple physical addresses 176
- multiple port device 179, 185, 226
- multiple recovered errors 195
- multiple target 11
- multiple-bit field 193
- multiple-byte field 193
- multiple-initiator system 201, 202
- Multiplication Factor
 - Caching Parameters page 134, 135

N

- NACA. see Normal Auto Contingent Allegiance
- ND. see Notched Drive
- negation period 15
- new microcode 216
- nexus 5, 7
- no additional sense information 193
- no defect spare location available 177
- no operation 32
- no sense 169
- no sense sense key 195
- nonarbitrating system 18, 24
- Non-Cache Segment Size
 - Caching Parameters page 134, 136
- non-changeable parameter 113, 117, 119, 149
- non-changeable value 112, 113, 116, 124, 127, 129
- nonfatal error 191
- non-medium error page 108
- non-recoverable error 6
- non-recoverable hardware failure 195
- nonrecovered error 195
- non-supported page code 122
- non-volatile memory 84, 87, 105, 217
- nonzero value 113
- NormACA. see Normal ACA Supported
- Normal ACA Supported
 - Drive Standard Inquiry Data 88, 89
- normal auto contingent allegiance 46, 48
- normal auto contingent allegiance bit 67
- normal phase 199
- not ready 209, 211
 - sense key 195
- notch 128
- notch page 115, 120, 139, 150
 - descriptor 122
- Notched Drive
 - Notch page 139
- Number of Blocks
 - Mode Select (10) parameter list 117
 - Mode Select parameter list 113
 - Mode Sense Data 121
 - Mode Sense Parameter List 151
 - Synchronize Cache command 210

- Write Same command 222
- number of blocks 151
- Number of Cache Segments
 - Caching Parameters page 134, 136
- Number of Cylinders
 - Rigid Drive Geometry Parameters page 130, 131
- Number of Heads
 - Rigid Drive Geometry Parameters page 130, 131
- Number of Source Descriptors
 - Rebuild and Regenerate parameter data 179

O

- object 5
- Obsolete
 - Read Reservations Descriptor 156
- obsolete 8
- Offset Boundary
 - Read Buffer Descriptor mode 165, 166
- offset boundary 218
- operating definition 72, 73
 - current 93
 - default 93
- operation
 - I/O 5
- Operation Code
 - Log Sense command 110
 - Persistent Reserve In command 153
 - Persistent Reserve Out command 158
 - Prefetch command 161
 - Read Defect Data (12) command 171
 - Rebuild command 178
 - Regenerate command 185
 - Release (10) command 188
 - Report LUNs command 189
 - Reserve (10) command 202
 - Write Same command 222
 - XDRead command 223
 - XDWrite command 225
 - XDWrite Extended command 226
 - XPWrite command 228
- operation code 43, 44, 76, 77
- option 8
 - implementation 4
 - protocol 6
- optional page format 181
- ordered task 33
- ordered task queue tag message 71, 72
- ordered task set tag 32
- OR-tied signals 13
- output
 - SCSI bus signal 12
- overlapped command 67
- overlapped commands attempted 67, 73
- overlapped tasks attempted 33

P

P bit 169, 170, 171, 172

P list 171

P type defect 81

packetized command structure 43

Pad

Rebuild and Regenerate parameter data 179, 180

Pad Length

Rebuild and Regenerate parameter data 179

Page Code

Caching Parameters page 134

Control Mode page 137

Device Behavior page 101

Device Identification page 94

Disconnect/Reconnect Control page 126

Error Recovery page 123

ETF Log Date Code page 98

Firmware Numbers page 96, 97

Format Parameters page 128

Implemented Operating Definition page 92, 93

Informational Exceptions Control page 146

Inquiry command 87

Jumper Settings page 99

Log Page format 104

Log Sense command 110

Mode Select page descriptor header 114

Mode Sense (10) command 149, 150

Mode Sense (6) command 119, 120

Mode Sense page descriptor header 122

Notch page 139

Power Condition page 141

Rigid Drive Geometry Parameters page 130

Supported Diagnostic page 207

Supported Diagnostic Pages 183

Supported Log page 111

Supported Vital Product pages 91

Translate Address page 207

Translate Address page—Receive Diagnostic 183

Unit Attention Parameters page 144

Unit Serial Number page 92

Verify Error Recovery page 132

Xor Control Mode page 142

page code 114, 122

Page Control

Log Select command 102, 103

Log Sense command 110

Page Control Field

Mode Sense (10) command 149

Mode Sense (6) command 119

page control field 149

bit 119

page descriptor 117

Page Descriptor Header

Caching Parameters page 134

Page Format

Mode Select (10) command 116

Mode Select (6) command 112

Send Diagnostic command 206

page format

optional 181

page header 151

Page Length

Caching Parameters page 134

Control Mode page 137

Device Behavior page 101

Device Identification page 94

Disconnect/Reconnect Control page 126

Error Recovery page 123

ETF Log Date Code page 98

Firmware Numbers page 96, 97

Format Parameters page 128

Implemented Operating Definition page 92

Informational Exceptions Control page 146

Jumper Settings page 99

Log Page format 104

Mode Select page descriptor header 114

Mode Sense page descriptor header 122

Notch page 139

Power Condition page 141

Rigid Drive Geometry Parameters page 130

Supported Diagnostic page 207

Supported Diagnostic Pages 183

Supported Log page 111

Supported Vital Product pages 91

Translate Address page 207

Translate Address page—Receive Diagnostic 183

Unit Attention Parameters page 144

Unit Serial Number page 92

Xor Control Mode page 142

page length 119, 122

Pages Notched

Notch page 139, 140

parameter

changeable 113, 117, 149

invalid 43

low-level hardware 72

mode 144

non-changeable 113, 117, 119, 149

ratio 126

rounded 144

savable 112

unchangeable 144

Parameter Code

Cache Statistics page 108

Error Counter pages 107

Log Parameter 105

Non-Medium Error page 108

Parameter Code Reset

Log Select command 102

parameter control byte 105

parameter counter 105

Parameter Data Length

- Change Definition command 78
- Parameter Information
 - Mode Select (10) parameter list 117
 - Mode Sense Data 121
- Parameter Length
 - Log Parameter 105, 106
- parameter list
 - Persistent Reserve Out 159
- Parameter List Length
 - Log Select command 102, 103
 - Mode Select (10) command 116
 - Mode Select (6) command 112
 - Persistent Reserve Out command 158
 - Rebuild command 178, 179
 - Regenerate command 185, 186
 - Send Diagnostic command 206
- parameter list length 46, 216, 217, 218
 - error 158
- Parameter Pointer
 - Log Sense command 110, 111
- Parameter Pointer Control
 - Log Sense command 110
- Parameter Savable
 - Caching Parameters page 134
 - Control Mode page 137
 - Disconnect/Reconnect Control page 126
 - Error Recovery page 123
 - Informational Exceptions Control page 146
 - Mode Sense page descriptor header 122
 - Notch page 139
 - Power Condition page 141
 - Rigid Drive Geometry Parameters page 131
 - Unit Attention Parameters page 144, 145
 - Verify Error Recovery page 132, 133
- Parameter Value
 - Log Parameter 105, 106
- parameter value
 - rounding 72
- Parameters Savable
 - Format Parameters page 128
 - Rigid Drive Geometry Parameters page 130
- parameters savable 122
- parity
 - data 12
- Parity Disable 99
- Parity Enable
 - Jumper Settings page 99
- parity error 41
- Partial Medium Indicator
 - Read Capacity command 167
- partitions 1
- path control 26
- Pattern Type
 - Initialization Pattern Descriptor 84
- PBdata. see Physical Block Data
- PC. see Page Control
- PCF. see Page Control Field
- PCR. see Parameter Code Reset
- PE. see Parity Enable
- peer entities 5
- peer-to-peer protocol service 5
- pending task 5
- PER. see Post Error
- Perf. see Performance
- Performance
 - Informational Exceptions Control page 146
- Peripheral Device Type
 - Device Behavior page 101
 - Device Identification page 94
 - Drive Standard Inquiry Data 88
 - ETF Log Date Code page 98
 - Firmware Numbers page 96, 97
 - Implemented Operating Definition page 92
 - Jumper Settings page 99
 - Supported Vital Product pages 91
 - Unit Serial Number page 92
- Peripheral Qualifier
 - Device Behavior page 101
 - Device Identification page 94
 - Drive Standard Inquiry Data 88
 - ETF Log Date Code page 98
 - Firmware Numbers page 96, 97
 - Implemented Operating Definition page 92
 - Jumper Settings page 99
 - Supported Vital Product pages 91
 - Unit Serial Number page 92
- persistent reservation scope code 156
- Persistent Reserve In command 48, 153, 158
- persistent reserve in parameter data 154, 155
- persistent reserve out
 - parameter list 159
 - service action 158, 160
 - valid parameters 160
- Persistent Reserve Out command 48, 67, 153, 154, 155, 158
- PF. see Page Format
- phase
 - command 22
 - data 22
 - data in 22
 - data out 22
 - message 23
 - message in 23
 - message out 23
 - reselection 19
 - sequence, SCSI bus 24
 - status 23
- phase sequence 25
- physical address 222
 - multiple 176
- Physical Block Data
 - Write Same command 222
- physical interconnect 6
- physical interface
 - characteristics 8
- physical medium 210

- Physical Sector
 - Defect List Header Description 170
 - Read Defect Data (12) list header 172
 - physical sector 129
 - physical sector address 183, 207
 - format 184
 - physical sector format 82, 85, 86, 169, 171, 222
 - PMI. see Partial Medium Indicator
 - pointer
 - current 26
 - saved 26
 - port 6, 227
 - Port Control
 - Rebuild command 178, 179
 - Regenerate command 185
 - XDWrite Extended command 226, 227
 - Post Error
 - Error Recovery page 123, 124
 - Verify Error Recovery page 132, 133
 - post error bit 148
 - power condition page 115, 120, 141, 150
 - descriptor 122
 - power consumption 209
 - power cycle 144, 147, 216, 217
 - power dissipation 145
 - power on cycle 218
 - power on reset 119, 149
 - power on time parameter 109
 - power on, reset, or bus device reset occurred 216
 - power up test error 199
 - PPC. see Parameter Pointer Control
 - predicted failure 147
 - Preempt
 - Service Action code 159
 - preempt 48
 - Preempt and Abort
 - Service Action code 159
 - preempt and abort service action 154
 - preempt and clear 48
 - preempt and clear action 67
 - preempt and clear service action 155, 160
 - preempt service action 154, 155, 160
 - prefetch 80
 - prefetch and multisegmented cache control 162
 - Prefetch command 161
 - Prevent/Allow Medium Removal command 161
 - previous reservation state 201
 - primary defect 81
 - procedure 6
 - reselection 20
 - SCSI 6
 - process
 - I/O 17
 - Product Identification
 - Drive Standard Inquiry Data 88, 90
 - Product Revision Level
 - Drive Standard Inquiry Data 88, 90
 - Product Serial Number
 - Unit Serial Number page 92
 - product serial number 95
 - Product Unique Sense Data
 - Extended Sense Data Format 192
 - Progress Indication
 - Format Progress Indication Bytes 194
 - protocol 6
 - general message 27
 - layer 6
 - lower level 5
 - option 6
 - SCAM 19
 - SCSI 6
 - SCSI application 6
 - SCSI-3 standards 2
 - service confirmation 6
 - service indication 6
 - service request 6
 - service response 6
 - service, unconfirmed 7
 - upper level 8
 - protocol service
 - confirmed 4
 - peer-to-peer 5
 - protocol-specific 8
 - PS. see Parameter Savable
- ## Q
- QErr. see Queue Error Management
 - queue 6, 138
 - queue a command 47
 - Queue Algorithm Modifier
 - Control Mode page 137, 138
 - Queue Error Management
 - Control Mode page 137, 138
 - queue full status 48
 - queue tag 6
 - queue tag value 72
 - queued environment 41
 - queued task 71
 - quick arbitration 25
- ## R
- RAC. see Report A Check
 - RAREA. see Reserved Area
 - ratio parameter 126
 - RC. see Read Continuous
 - RCD. see Read Cache Disable
 - Read Buffer command 164, 166, 215
 - read buffer descriptor 166, 217, 218
 - mode 165
 - read buffer header 165, 216
 - Read Cache Disable
 - Caching Parameters page 134, 135
 - read cache disable bit 164, 178, 185, 219, 225, 226, 228
 - Read Capacity command 45, 72, 167, 168

- read capacity data 168, 204
- read combined descriptor header and data 164, 165
- Read command 49, 50, 80, 143, 162, 173, 178, 185, 213
- Read Continuous
 - Error Recovery page 123, 124
- read data 164, 165
- Read Data Buffer
 - Write Buffer command 215
- read data buffer 209
- read data from echo buffer 164, 166
- Read Defect Data (10) command 169
- Read Defect Data (12) command 171
- read defect data (12) list header 172
- read defect data transfer 169
- read descriptor 164
- Read Element Status command 172
- read error
 - unrecoverable 177
- read error counter pages 107
- Read Exclusive
 - Persistent Reservation 157
- Read Extended command 173
- Read Keys
 - Service Action code 154
- Read Long command 175, 221
- read look ahead 80
- read operation 174
- read reservation service action 156
- Read Reservations
 - Service Action code 154
- Read Retry Count
 - Error Recovery page 123, 124
- Read Shared
 - Persistent Reservation 157
- read transfer 224
- read/write heads 131
- read-write error recovery page 175
- Reassign Blocks command 176
- reassign blocks defect list 177
- rebuild 224
- rebuild and regenerate parameter data 179
- Rebuild command 178
- Rebuild Delay
 - Xor Control Mode page 142, 143
- Rebuild Length
 - Rebuild command 178, 179
- rebuild length 185
- rebuild parameter list 178, 185
- rebuild size 142
- rebuild/regenerate length 180
- Receive Diagnostic
 - Translate Address page 183
- receive diagnostic results 209
- Receive Diagnostic Results command 181, 183, 206, 207
- receiver 6
- reconnect 127, 213
- recoverable error event 108
- recoverable read error 174
- recovered error 124, 133, 148, 162, 163, 169, 171, 174, 193, 227
- recovered error sense key 72, 195
- recovered read error 163
- recovery level 124
- Recovery Time Limit
 - Error Recovery page 123, 125
- recovery time limit 133
- redundancy group 223
- reference model 6
- regenerate 224
- Regenerate command 185, 186, 223
- Regenerate Length
 - Regenerate command 185, 186
- regenerate length 185
- Register
 - Service Action code 159
- register service action 154, 155, 157, 159, 160
- registered reservation key 160
- RelAdr. see Relative Address
- Relative Address
 - Drive Standard Inquiry Data 88, 90
 - Prefetch command 161
 - Read Capacity command 167
 - Read Extended command 173
 - Read Long command 175
 - Synchronize Cache command 210
 - Verify command 212
 - Write and Verify command 214
 - Write Extended command 219
 - Write Long command 221
 - Write Same command 222
- relative address 45, 174
 - bit 219
- Release
 - Service Action code 159
- release 187, 209
- Release (10) command 188
- Release (6) command 187
- Release command 48, 72, 158, 187, 188, 200, 202
- release service action 154, 155, 159, 160
- Report A Check
 - Control Mode page 137, 138
- Report Count
 - Informational Exceptions Control page 146, 147
- Report Log Exception Condition
 - Control Mode page 137
- report log exception condition 105, 106
- Report LUNs command 189
- report LUNs parameter list format 190
- reporting information exceptions 148
- REQ. see Request
- REQ/ACK
 - handshake 21, 22, 23, 32
 - offset 22, 35, 36

- request 6, 7
 - confirmation transaction 6
 - response transaction 6
 - SCSI bus signal 12
 - request sense 20, 45, 46, 144, 209
 - Request Sense command 43, 67, 68, 70, 102, 175, 191, 192, 199, 200, 221
 - rescind reservation 201
 - reselection 127
 - reselection phase 19
 - reselection procedure 20
 - resend message 41
 - reservation 165, 187, 200, 201
 - reservation access conflict 213, 219
 - reservation conflict 157
 - reservation conflict status 48, 153, 159, 162, 171, 200, 213, 219
 - reservation descriptor 156
 - Reservation Identification
 - Release (10) command 188
 - Release (6) command 187
 - Reserve (10) command 202
 - Reserve (6) command 200
 - reservation identification byte 187
 - Reservation Key
 - Persistent Reserve Out parameter list 159
 - Read Reservations Descriptor 156
 - reservation key 154, 159, 160
 - Reservation Key List
 - Persistent Reserve In parameter data 154, 155
 - reservation process 201
 - reservation protocol 201
 - Reservations Descriptor
 - Persistent Reserve In parameter data 155
 - Reserve
 - Service Action code 159
 - reserve 209
 - Reserve (10) command 202
 - Reserve (6) command 200
 - Reserve command 48, 72, 158, 200, 201
 - reserve service action 154, 155, 159, 160
 - reserved 8
 - Reserved Area
 - Translate Address page—Receive Diagnostic 183, 184
 - reserved bit 43
 - reserved byte 43
 - reserved code value 43
 - reserved fields 43
 - reserved nonuser-accessible cylinder 206
 - reset 144, 217
 - command 213
 - condition 23, 24, 162
 - event 6
 - hard 4
 - hold time 15
 - message 144
 - SCSI bus signal 12
 - response 6
 - service 7
 - Response Data Format
 - Drive Standard Inquiry Data 88, 89
 - restore pointers 31, 33, 41
 - restriction
 - signal 23
 - retry 124, 133
 - retry algorithm 108, 124
 - retry count field 194
 - retry mechanisms 133
 - retry message 41
 - retry phase 41
 - return all supported pages 115, 120, 150
 - Return Changeable values
 - Page Control Field bits 119
 - Return Current values
 - Page Control Field bits 119
 - Return Default values
 - Page Control Field bits 119
 - return defect list header 169, 170
 - Return Saved values
 - Page Control Field bits 119
 - Rezero Unit command 203
 - rigid drive geometry page descriptor 122
 - rigid drive geometry parameters 115, 120, 150
 - page 130
 - RLEC. see Report Log Exception Condition
 - RMB
 - Format Parameters page 128
 - ROM release date 96
 - Rotational Offset
 - Rigid Drive Geometry Parameters page 130, 131
 - Rotational Position Locking
 - Rigid Drive Geometry Parameters page 130, 131
 - rotational position locking 131
 - Round
 - Unit Attention Parameters page 144
 - round down 138
 - rounded parameter 72, 144
 - rounding 113, 117
 - rounding of parameter values 72
 - RPL. see Rotational Position Locking
 - RST. see Reset
- ## S
- S.M.A.R.T. see Self-Monitoring Analysis and Reporting Technology
 - SAM-compliant target identifier 180, 226
 - SAP. see Servo Adaptive Parameters
 - savable parameter 112
 - Save
 - Change Definition command 78
 - Save Control
 - Change Definition command 78

- save data pointer 33, 41
- save data pointer message 50
- save executable microcode 215
- Save Implemented
 - Implemented Operating Definition page 92, 93
- save log counter 203
- Save Mode Parameters
 - Mode Select (10) command 116
 - Mode Select (6) command 112
- save mode parameters 122
- save mode parameters bit 119, 131, 149
- save mode parameters element reservation 153
- save parameter 105
- Save Parameters
 - Log Select command 102
 - Log Sense command 110
- saved log parameter 203
- saved pointer 26
- saved value 119, 149
- SAVIMP. see Save Implemented
- SCAM
 - compliant 18
 - devices 19
 - initiator 19
 - Product Manual 19, 233
 - protocol 19, 233
 - requirements 19
 - selection 19
 - target 19
 - tolerant 19
- SCC Supported
 - Drive Standard Inquiry Data 88, 89
- SCCS. see SCC Supported
- Scope
 - Persistent Reserve Out command 158
 - Read Reservations Descriptor 156
- Scope-specific Address
 - Persistent Reserve Out parameter list 159, 160
 - Read Reservations Descriptor 156
- SCSI
 - dual-port 40
- SCSI application
 - layer 6
 - protocol 6
 - protocol transaction 6
- SCSI architecture model
 - SCSI-3 standards 2
- SCSI bus 11, 49, 50
 - condition 194
 - conditions and miscellaneous features supported 131
 - ID 187, 188, 200
 - phase sequence 24
 - procedure 43
 - signals 12
 - timing 13
- SCSI command
 - structure 43
 - summary 9
- SCSI compliance level 72
- SCSI configured automatically. see SCAM
- SCSI device 6, 233
 - identifier 6
 - reservations 70
- SCSI Firmware Release Number
 - Firmware Numbers page 96, 97
- SCSI I/O system 6, 7
- SCSI ID bits 11
- SCSI initiator 48
- SCSI interface parity error 195
- SCSI message
 - handling 40
 - summary 9, 10
- SCSI port 155
 - control mode page 115, 120
- SCSI protocol 43
 - layer 6
- SCSI specification level 72
- SCSI-1 89, 125
 - device 79
- SCSI-2 144
 - device 79
 - feature 137
 - Unit Attention Parameters page 144, 145
- SCSI-3 145
 - architectural model 148
 - standards 1
- SDTR. see Synchronous Data Transfer Request
- SE. see Single-ended
- Search Data command 195
- Search Data Equal command 203
- Search Data High command 203
- Search Data Low command 203
- Secondary Address
 - XDWrite Extended command 226
- Secondary Logical Block Address
 - XDWrite Extended command 226
- secondary logical block address 227
- secondary target identifier 226
- sector size 129
- Sectors per Track
 - Format Parameters page 128, 129
- Seek command 204, 205
- seek error 203
- Seek Extended command 205
- Segment Number
 - Extended Sense Data Format 192
- SEL. see Select
- select
 - SCSI bus signal 12
- selection abort time 15
- selection phase 18
- selection timeout delay 15
- Self Seek
 - Unit Attention Parameters page 144, 145
- Self Test

- Send Diagnostic command 206
- self-describing microcode 217
- self-monitoring analysis and reporting technology
 - 42, 146
 - measurement 109
 - system 42, 146
- Send Diagnostic command 181, 183, 184, 206, 207
- sender 6
- sense code to overlapped commands attempted 72
- sense data 41, 67, 68, 142, 162, 177, 178, 179, 185, 191, 199, 223, 226
- sense data error 216
- Sense Data Length
 - Mode Sense Data 121
 - Mode Sense Parameter List 151
- Sense Key
 - Extended Sense Data Format 192
- sense key 41, 43, 68, 73, 79, 83, 85, 103, 104, 106, 110, 116, 122, 124, 127, 133, 144, 148, 158, 160, 161, 163, 169, 171, 174, 175, 177, 189, 192, 193, 194, 195, 200, 206, 207, 209, 210, 211, 212, 213, 214, 216, 217, 218, 220, 221, 227
 - descriptions 195
- Sense Key Specific
 - Extended Sense Data Format 192, 193
- sense key specific field 193, 194
- Sense Key Specific Valid
 - Extended Sense Data Format 192, 193
- sense key to aborted command 72
- sequencing of commands 66
- sequential track 129
- server 6
- service 6
 - delivery failure 6
 - delivery port 7
 - delivery subsystem 7
 - delivery transaction 7
- Service Action
 - Persistent Reserve In command 153
 - Persistent Reserve Out 158, 159
 - Persistent Reserve Out command 158
- service action 154, 155, 160
- Service Action Reservation Key
 - Persistent Reserve Out parameter list 159, 160
- service response 7
- Servo Adaptive Parameters Firmware Release Number
 - Firmware Numbers page 97
- Servo Adaptive Parameters Manufacturing Key
 - Firmware Numbers page 98
- Servo Firmware Product Family and Product Family Member IDs
 - Firmware Numbers page 98
- Servo RAM Release Date In ASCII
 - Firmware Numbers page 97
- Servo RAM Release Number
 - Firmware Numbers page 97
- Servo ROM Release Date In ASCII
 - Firmware Numbers page 97
- Servo ROM Release Number
 - Firmware Numbers page 96, 97
- Set Limits command 208
- shall 8
- Shared Access
 - Persistent Reservation 157
- should 8
- signal 7
 - OR-tied 13
 - sources 13
 - values 13
- signal restrictions 23
- simple queue tag message 138
- simple task 33
- simple task queue tag message 71
- simple task tag 32
- single command example 49, 50
- single initiator 11
- single target 11
- single-ended 12
- Size Enable
 - Caching Parameters page 134, 135
- SIZE. see Size Enable
- skew 129
- SKSV. see Sense Key Specific Valid
- SMP. see Save Mode Parameters
- Soft Reset 88
 - Drive Standard Inquiry Data 90
- Soft Sectoring
 - Format Parameters page 128
- Source Descriptor
 - Rebuild and Regenerate parameter data 179
- Source Descriptor format
 - Rebuild and Regenerate parameter data 180
- Source Descriptor/Pad Length
 - Rebuild and Regenerate parameter data 179
- source device 7, 178
- Source Device Address
 - Source Descriptor format 180
- Source Starting Logical Block Address
 - Source Descriptor format 180
- SP. see Save Parameters
- spare sector 129
- spare track 129
- spindle rotation 209
- spindle synchronization 131
- SSEC. see Soft Sectoring
- SSM. see Synchronous Select Mode
- stable error correction code syndrome 108
- standard inquiry data 95, 144
 - format 88, 189
- Standby
 - Power Condition page 141
- Standby Condition Timer

- Power Condition page 141
- Start
 - Start/Stop Unit command 209
- Start command 211
- start unit 199
- Start/Stop Unit command 209
- Starting Boundary
 - Notch page 139, 140
- Starting Cylinder—Reduced Write Current
 - Rigid Drive Geometry Parameters page 130
- Starting Cylinder—Write Precomp
 - Rigid Drive Geometry Parameters page 130
- status 209
- status byte 43, 47
- status phase 23, 41, 45, 47, 162, 213
- Stop Format
 - Defect List header 83
- storage array 223
- storage device 223
- storage medium 157
- stored pointer 49
- STPF. see Stop Format
- Strict
 - Unit Attention Parameters page 144
- strict 116, 117, 144
- strict mode 113
- subsystem 7
- superseding reservations 201
- superseding reserve command 201
- supervised Xor operation 224
- Supplied Format
 - Translate Address page 207
 - Translate Address page—Receive Diagnostic 183
- support services 229
- Supported Diagnostic page
 - Send Diagnostic command 207
- supported diagnostic pages 183
- Supported Log page
 - Log Sense command 111
- Supported Operating Definition List
 - Implemented Operating Definition page 92
- Supported Page List
 - Supported Diagnostic Pages 183
 - Supported Log page 111
 - Supported Vital Product pages 91
- SURF
 - Format Parameters page 128
- suspended information 7
- Sync. see Synchronous Transfer
- synchronization command 199
- synchronize cache 210
- Synchronize Cache command 210
- synchronized spindle
 - master 131
 - slave 131
- synchronizing process 199
- synchronous data transfer 22

- synchronous data transfer request 33, 34, 35, 36, 144
 - message 34
 - response 36
- synchronous negotiation 41
- Synchronous Select Mode
 - Unit Attention Parameters page 144
- Synchronous Transfer
 - Drive Standard Inquiry Data 88, 90

T

- T10/1302D 25
- Table Address
 - XDWrite Extended command 226
- tag 7, 32
- tag value 67
- tagged command 71
- tagged I/O command 72
- tagged overlapped tasks 67
- tagged queuing 137, 138
- tagged task 67
 - address 7
 - queuing 32, 71
- target 2, 7
 - multiple 11
 - single 11
- target hard reset 70
- target identifier 226
- target reset 163, 174, 213, 220
- target reset task management request 70
- Target Save Disable
 - Log Parameter 105
- target save disable 102
- target transfer disable 34
- task 7, 66, 67, 165
 - abort event 7
 - address 7, 67
 - and command lifetimes 67
 - attribute 30
 - attribute message 32, 33
 - complete 7, 30, 66
 - completed 3
 - completion event 7
 - current 4
 - ended 7
 - ended event 7
 - faulting 4
 - management 32
 - management function 7, 48
 - management request 7
 - management response 7
 - manager 7
 - pending 5
 - queue tag message 71, 72
 - queued 71
 - slot 7
 - tag 7

- task set 7, 47, 48
 - faulted 4
 - full status 48
 - management 66, 67
 - queue 33
 - task state
 - dormant 4
 - enabled 4
 - TB. see Transfer Block
 - TDD. see Target Transfer Disable
 - technical support services 229
 - temporary initiator 178
 - terminate 169
 - terminate command 113, 117, 169, 195, 206, 212
 - Terminate Task
 - Drive Standard Inquiry Data 88, 89
 - termination 8
 - Terminator Enable
 - Jumper Settings page 99, 100
 - Terminator Power
 - Jumper Settings page 100
 - Test
 - Informational Exceptions Control page 146, 147
 - test unit ready 209
 - Test Unit Ready command 72, 211
 - thermal calibration 102, 105
 - thermal compensation 203
 - Third Party
 - Release (10) command 188
 - Release (6) command 187
 - Reserve (10) command 202
 - Reserve (6) command 200
 - third party
 - bit 201
 - command 7
 - ID field 187, 201
 - release 187
 - reservation 187, 201, 202
 - Third Party Device ID
 - Release (10) command 188
 - Release (6) command 187
 - Reserve (10) command 202
 - Reserve (6) command 200
 - Threshold Met Criteria
 - Log Parameter 105, 106
 - timing
 - examples 51
 - SCSI bus 13
 - waveforms 51
 - TMC. see Threshold Met Criteria
 - total (re-writes or re-reads) code 107
 - total bytes processed code 108
 - total error corrected code 107
 - total times correction algorithm processed code 108
 - total uncorrected errors code 108
 - Track Skew Factor
 - Format Parameters page 128, 129
 - Tracks per Zone
 - Format Parameters page 128, 129
 - transaction 7
 - Transfer Block
 - Error Recovery page 123, 124
 - Transfer Disable
 - Drive Standard Inquiry Data 88, 90
 - transfer disable bit 30, 34
 - Transfer Length
 - Prefetch command 161
 - Read command 162
 - Read Extended command 173, 174
 - Write and Verify command 214
 - Write command 213
 - Write Extended command 219
 - XDRead command 223
 - XDWrite command 225
 - XDWrite Extended command 226, 227
 - XPWrite command 228
 - transfer length 46, 173, 185, 216, 217, 225
 - transfer option
 - fast synchronous 15
 - transfer period 15
 - factor field 35
 - transfer sense data 191
 - transfer width exponent 37
 - Translate Address
 - Translate Address page—Receive Diagnostic 183
 - translate address operation 206
 - Translate Address page 183
 - Receive Diagnostic 183
 - Send Diagnostic command 207
 - Translate Format
 - Translate Address page 207
 - Translated Address
 - Translate Address page—Receive Diagnostic 183, 184
 - translated address 184, 207
 - Translated Format
 - Translate Address page—Receive Diagnostic 183, 184
 - TrmTsk. see Terminate Task
 - TrnDis. see Transfer Disable
 - TSD. see Target Save Disable
 - TST 66
 - Type
 - Persistent Reserve Out command 158
 - Read Reservations Descriptor 156
- ## U
- unassigned ID 19
 - unchangeable parameter 144
 - unconfirmed protocol service 7
 - unexpected bus free 41, 47
 - condition 199
 - unexpected disconnect 31
 - Unit Attention

- Unit Attention Parameters page 144
- unit attention 144, 163, 174, 220
- unit attention condition 69, 70, 73, 102, 105, 112, 138, 216, 217, 218
- unit attention mode parameters page 70
- unit attention page parameters 115, 120, 150
- unit attention parameters 69
 - page 144
 - page descriptor 122
- unit attention sense key 70, 195
- Unit Off Line
 - Send Diagnostic command 206
- unit serial number 95
 - page 91, 92
- UnitOFL. see Unit Off Line
- unlinked command 8
- unlock of overlapped area 210
- unrecoverable error 108, 124, 133, 162
- unrecoverable read error 163, 177
- unrecovered error 133, 199
- unrecovered read error 174
- unsuccessful I/O process termination 17
- unsupported bit 119
- unsupported field 119, 149
- unsupported page code 113, 117
- unsupported value 113
- untagged command 33, 72
 - environment 127
- untagged task 67
 - address 7
 - queuing 71
- Unused Vendor Specific Area
 - Drive Standard Inquiry Data 88
- update write 224
- upper level protocol 8

V

- valid bit 177
- valid initiator reconnection 30
- Validity Bit
 - Extended Sense Data Format 192
- value
 - changeable 126, 129, 134, 138, 149
 - current 149
 - default 149
 - interleave 80
 - non-changeable 112, 113, 116, 124, 127, 129
 - nonzero 113
 - saved 149
 - unsupported 113
- variable track capacity recording 128
- vendor ID 95, 96
- Vendor Identification
 - Drive Standard Inquiry Data 88, 90
- Vendor Specific
 - Defect List header 83, 84
 - Drive Standard Inquiry Data 88, 89

- vendor specific 8
 - condition 217
 - error recovery 67
 - microcode 216
 - sense key 195
- vendor unique code 182
- vendor unique executable microcode 216
- Vendor Unique Fault Information
 - Diagnostic Data Bytes 182
- Verification Length
 - Verify command 212
- verify 212
- Verify command 212
- Verify Correction Span
 - Verify Error Recovery page 132, 133
- verify error counter page 107
- verify error recovery page 132
 - descriptor 122
 - parameters 120, 150
- verify error recovery parameters 115
- verify function 83
- Verify Recovery Time Limit
 - Verify Error Recovery page 132, 133
- Verify Retry Count
 - Verify Error Recovery page 132, 133
- verify retry count 133
- Version Number
 - Device Behavior page 101
- virtual logical unit 93
- visible medium position 206
- vital product data page 91, 96
 - supported 97, 98
- volume overflow 220
 - sense key 195
- VS. see Vendor Specific

W

- WBus16. see Wide Bus 16
- WBus32. see Wide Bus 32
- WCE. see Write Cache Enable
- WDTR. see Wide Data Transfer Request
- Wide Bus 16
 - Drive Standard Inquiry Data 88, 90
- Wide Bus 32
 - Drive Standard Inquiry Data 88, 90
- wide data transfer 36
- wide data transfer capability 39
- wide data transfer request 34, 36, 37, 38, 144
 - response 37
- wide SCSI byte ordering 38
- WP. see Write Protect
- Write and Verify command 214
- Write Buffer command 165, 166, 215, 216, 217, 218
- write buffer header 216
- write cache 219
- Write Cache Enable

Caching Parameters page 134, 135
write cache enable bit 178, 185, 219, 225, 226,
228
Write command 127, 178, 185, 213, 219, 225, 226
write data buffer 209
Write Data Only mode
 Write Buffer command 216
Write Data to Echo Buffer mode
 Write Buffer command 218
write error counter pages 107
Write Exclusive
 Persistent Reservation 157
Write Extended command 219
Write Long command 221
write operation 218, 219
Write Protect
 Jumper Settings page 99
 Mode Sense Data 121
 Mode Sense Parameter List 151
Write Retention Priority
 Caching Parameters page 134, 135
Write Retry Count
 Error Recovery page 123, 124
Write Same command 222
write sequence 224
write transfer 224
write, read, and verify error counter pages 107
writes prohibited 157

X

XDRead command 223
XDWrite command 142, 225
XDWrite Extended command 142, 226, 227
Xor command 223, 225, 226, 228
Xor control mode page 115, 120, 142, 150, 223
 descriptor 122
Xor data 178, 185, 227
Xor Disable
 Xor Control Mode page 142
Xor operation 178, 180, 224
Xor result data 179, 228
XorDis, see Xor Disable
XPWrite command 142, 226, 227, 228
XPWrite target 227

Z

ZBR zone 128, 129



Seagate Technology, Inc.
920 Disc Drive, Scotts Valley, California 95066-4544, USA
Publication Number: 77738479, Rev. J, Printed in USA