Hard disk drive specifications

# Ultrastar 36Z15

3.5 inch SCSI hard disk drive

Models: IC35L036UWPR15
IC35L036UCPR15
IC35L018UWPR15
IC35L018UCPR15

Revision 1.0

10 April 2001

Publication #3625

This page intentionally left blank.

*IBM STORAGE PRODUCTS*

# Hard disk drive specifications

# Ultrastar 36Z15

3.5 inch SCSI hard disk drive

Models: IC35L036UWPR15
IC35L036UCPR15
IC35L018UWPR15
IC35L018UCPR15

---

Revision 1.0 10 April 2001

Publication #3625

---

**1st Edition (Rev. 0.1) January 18, 2001 (Preliminary)**
**2nd Edition (Rev. 0.2) January 30, 2001 (Preliminary)**
**3rd Edition (Rev. 0.3) January 31, 2001 (Preliminary**
**4th Edition (Rev. 0.4) February 1, 2001 (Preliminary)**
**5th Edition (Rev 0.5) February 14, 2001 (Preliminary)**
**6th Edition (Rev 0.6) February 15, 2001 (Preliminary)**
**7th Edition (Rev 0.7) February 27, 2001 (Preliminary)**
**8th Edition (Rev 0.8) March 2, 2001 (Preliminary)**
**9th Edition (Rev 0.9) March 5, 2001 (Preliminary)**
**10th Edition (Rev 0.95) March 24, 2001 (Preliminary)**
**11th Edition (Rev 1.0) April 10, 2001**

# Table of Contents

# List of figures

# 1.0 General

This document contains the hardware and functional specifications for the Ultrastar 36Z15 family of high-capacity, 15,000-RPM, 3.5-inch disk drives. This drive offers formatted capacities of 36.7 GB and 18.4 GB in a one-inch high form factor with SCSI interface following the ANSI version 3 standard.

## 1.1 Introduction

The following model numbers reference this Ultra SCSI-160 version:

- IC35L036UCPR15-X (36GB, 80 pin SCSI 160)
- IC35L036UWPR15-X (36GB, 68 pin SCSI 160)
- IC35L018UCPR15-X (18GB, 80 pin, SCSI 160)
- IC35L018UWPR15-X (18GB, 68 pin, SCSI 160)

## 1.2 Glossary

| Word | Meaning |
|---|---|
| Kbpi | 1,000 bits per inch |
| Mbps | 1,000,000 bits per second |
| GB | 1,000,000,000 bytes |
| MB | 1,000,000 bytes |
| KB | 1,000 bytes, unless otherwise specified |
| Mb/sq.in | 1,000,000 bits per square inch |
| MLC | Machine Level Control |
| PFA | Predictive Failure Analysis (Trademark of IBM Corp.) |
| S.M.A.R.T. | Self-Monitoring Analysis and Reporting Technology |
| ADM | Automatic Drive Maintenance |
| SCAM | SCSI Configured Automatically |
| SE | Single-ended SCSI |
| LVD | Low Voltage Differential SCSI |
| FC-AL | Fibre Channel - Arbitrated Loop |

## 1.3 General caution

- User has the responsibility to adhere to all conditions specified in this document.
- User has the responsibility to mount the drive securely to minimize excess vibration, shock, and exposure to other unstable conditions.
- The breather hole located on top cover must not be obstructed.
- The drive can be damaged by Electric Static Discharge (ESD). Damage incurred to the drive after removing it from the shipping package and the ESD protective bag is the user's responsibility.

# 2.0 General features

- Storage capacities of 36.7 GB and 18.4 GB
- Both capacities are available with the following interface types:
  - Wide Ultra SCSI, Fast 40, LVD/SE multi-mode, without active termination, 68-pin connection
  - Wide Ultra SCSI, Fast 40, LVD/SE multi-mode, without active termination, 80-pin connection
  - Ultra3 SCSI, Ultra160, 160 MB/s, 68-pin connection
  - Ultra3 SCSI, Ultra160, 160 MB/s, 80-pin connection
- Interleave factor 1:1
- Variable sector sizing of 512 to 528 in 2-byte increments
- Tagged Command Queuing support (maximum queue depth of 128)
- Automatic Read/Write data transfer
- 4MB segmented sector buffer
  - Selectable to 128 KB x 27, 256 KB x13, or 512 KB x 6
- Average Read Seek time of 4.2 ms (36 GB model) and 3.4 ms (18 GB model)
- Adaptive Read Ahead algorithm
- Write Cache
- Back-to-back writes (merged writes)
- Five bytes per interval error correction (ECC) On-the-Fly
- Automatic defect reallocation
- Self diagnostics at power on
- Closed loop actuator servo
- Ramp Unload and Load for all heads providing no head/disk contact start stop
- Automatic inertia actuator latch
- Supports S.M.A.R.T.
- Glass substrate disks

This page intentionally left blank.

# Part 1. Functional specification

This page intentionally left blank.

# 3.0 Fixed Disk Subsystem Description

## 3.1 Control electronics

A microprocessor, logic modules, digital/analog modules, various drivers and receivers control the drive electrically. The control electronics perform the following major functions:

- Perform Basic Assurance Test Sequence (BATS) diagnostic self-check.
- Conduct a power on sequence and calibrate the servo.
- Monitor various timers, including those for head settling, servo failure, and track following.
- Analyze servo signals to provide closed loop control, including position error signal and estimated velocity.
- Monitor spindle rotational speed to control start, stop and rotating speed of the spindle.
- Control the voice coil motor (VCM) driver to align the actuator in a desired position.
- Monitor the actuator position and determine the target track for a Seek operation.
- Constantly monitor error conditions of the servo and take corresponding action if an error occurs.
- Control and interpret all interface signals between the host controller and the drive.
- Control Read/Write access of the disk media, including defect management and error recovery.

## 3.2 Head disk assembly

The head disk assembly (HDA) is assembled in a clean room and contains the disk stack, a spindle motor, actuator assembly, and VCM. Air in the HDA is constantly circulated and filtered once the drive is spun up. A breather filter located on the top cover ventilates the HDA.

A brushless, sensorless DC drive motor drives the spindle directly. Dynamic braking stops disk rotation quickly.

## 3.3 Actuator

The Read/Write heads are mounted in the actuator. The actuator is a rotary arm assembly driven by the VCM. A closed loop positioning servo system controls actuator movement. An embedded servo data pattern supplies feedback to the positioning servo to keep the Read/Write heads centered over the desired track.

The actuator assembly is balanced to allow vertical or horizontal mounting of the drive without any additional adjustment.

When the drive is powered off, the actuator automatically unloads the heads to a ramp located outside the disk area.

This page intentionally left blank.

# 4.0 Drive characteristics

## 4.1 Formatted capacity

| Description | 36-GB models | 18-GB models |
|---|---|---|
| Label capacity | 36.7 GB | 18.4 GB |
| Bytes per sector | 512 - 528 | 512 - 528 |
| Sectors per track | 322 to 465 | 372 to 465 |
| Number of heads | 12 | 8 |
| Number of disks | 6 | 4 |
| Data sectors per cylinder | 3864 to 5580 | 2480 to 3720 |
| Spare sectors per zone | 892 to 5652 | 892 to 5652 |
| Total data bytes per disk | 3,152,244,480 | 2,628,561,936 |
| Data cylinders per zone | 732 to 2213 | 732 to 2213 |
| Total logical data blocks | 71,687,340 (445DCAC) | 35,843,670 (222EE56) |
| Maximum addressable data blocks | 71,687,339 (445DCAB) | 35,843,669 (222EE55) |

Figure 1. Formatted capacity

## 4.2 Data sheet

| Buffer to/from media | 453 - 647 Mbps |
|---|---|
| Host to/from buffer (synchronous transfer rate):<br>• 68-pin/80-pin Fast 40<br>• 68-pin/80-pin Ultra160 | <br>80 MB/s<br>160 MB/s |
| Data buffer size | 4 MB |
| Number of buffer segments | 512KB x 6<br>256KB x 13<br>128KB x 27 |
| Rotational speed | 15,000 RPM |
| Recording density | 323 Kbpi min –<br>397 Kbpi max |
| Track density (average) | 27,000 TPI |
| Areal density (maximum) | 10.7 Gbits/in$^2$ |
| Data bands | 11/8 |

Figure 2. Data sheet

# 4.3 Performance characteristics

The following parameters characterize drive performance:

- Command overhead
- Mechanical head positioning
  - Seek time
  - Latency
- Data transfer speed
- Buffering operation (Read Ahead/Write Cache)
- Command sorting/Queuing Algorithms
- Interface Protocol handling/automation

**Note:** All of the above parameters contribute to drive performance. Other parameters contribute to the performance of the actual host system, where minimum and maximum values vary depending upon such factors as workload, logical and physical operating environments, and manufacturing process variations. This specification defines performance characteristics on the bare drive level so that system and application workloads can be approximated.

## 4.3.1 Command overhead

The table below shows minimum, maximum, and average command overhead in an unqueued environment for random single block transfers. Command overhead for a non-cache hit operation is the time extending from the last byte of the Command Data Byte (CDB) to Seek start.

The Time Measurement Window in an ICE (In Circuit Emulator) tool measured 1,000 commands with the following results:

| Description | Min (µs) | Avg (µs) | Max (µs) |
|---|---|---|---|
| Command receipt until Seek start | 48.96 | 52.48 | 111.96 |

Figure 3. Command receipt until Seek start

Command overhead for a cache hit is the time extending from the last byte of the CDB to the first byte of the data phase and excludes any initiator delay, including reconnections. The following values are used:

| **Read Command Case**<br>(Drive in quiescent state) | **Command Overhead Time**<br>(**µ**s) |
|---|---|
| Cache not hit | < 400 |
| Cache hit | < 21 |

Figure 4. Command overhead

## 4.3.2 Average Seek time (including settling)

| | Command type | Typical (ms) | Max (ms) |
|---|---|---|---|
| 36-GB model | Read | 4.2 | 4.8 |
| | Write | 4.7 | 5.3 |
| 18-GB model | Read | 3.4 | 3.7 |
| | Write | 3.9 | 4.2 |

Figure 5. Average Seek time

The performance specification uses "Typical" and "Max" throughout as follows:

**Typical**: Weighted average access time of the drive population tested at nominal environmental and voltage conditions.

**Max**: Maximum value measured on any one drive over the full range of the environmental and voltage conditions. (See Section 6.3, "Environmental requirements" on page 28 and Section 6.4, "D.C. power requirements" on page 31.)

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

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

*where*

```
Max   = Maximum Seek length
n     = Seek length (1 to max)
Tn.in = Inward measured Seek time for an "n" track Seek
Tn.out= Outward measured Seek time for an "n" track Seek
```

## 4.3.3 Full stroke Seek

| | Function | Typical (ms) | Max (ms) |
|---|---|---|---|
| 36-GB model | Read | 8.9 | 10.0 |
| | Write | 9.5 | 10.9 |
| 18-GB model | Read | 6.7 | 7.1 |
| | Write | 7.1 | 7.6 |

Figure 6. Full stroke Seek

Full stroke Seek is measured as the average of 1,000 full stroke Seeks with a random head switch from both inward and outward directions.

## 4.3.4 Head switch time (head skew)

| | |
|---|---|
| **Typical (ms)** | 0.84 |

Figure 7. Head switch time

Head switch time is the amount of time the fixed disk requires to access the next sequential block after reading the last sector in the current track.

## 4.3.5Track to track time

| | |
|---|---|
| **Typical (ms)** | 0.97 |

Figure 8. Track to track time

Track to track time refers to the amount of time the drive requires to access the next sequential block after completing a Read in the last sector of the current cylinder.

## 4.3.6 Average latency

| Rotation Speed | Time for one revolution (ms) | Average latency (ms) |
|---|---|---|
| 15,000 RPM | 4.0 | 2.0 |

Figure 9. Average latency

## 4.3.7 Drive ready time

| Model | Typical (sec) | Max (sec) |
|---|---|---|
| 36-GB model | 18.0 | 29.9 |
| 18-GB model | 14.0 | 29.9 |

Figure 10. Drive ready time

Ready time pertains to the state when the drive is immediately able to perform a media access command such as Read or Write. If the system receives a command during power on sequencing or before the drive reaches ready, drive ready time may extend beyond the time specified.

Power on includes the time required for the drive to complete internal self diagnostics.

## 4.3.8 Spindle stop time

| Model | Typical (sec) | Max (sec) |
|---|---|---|
| 36-GB model | 18.0 | 30 |
| 18-GB model | 14.0 | 30 |

Figure 11. Spindle stop time

The period from power off to the complete stop of the spindle is considered 'operating' state. Shock tolerance during this period is defined under 'operating shock'.

## 4.3.9 Data transfer speed

| **Disk - Buffer data transfer (typical)** | **36-GB model** (MB/s) | **18-GB model** (MB/s) |
|---|---|---|
| **zone 0** | | |
| Media transfer rate | 61.4 | 61.4 |
| Sustained disk to buffer | 52.8 | 52.8 |
| **ID zones** | | |
| Media transfer rate | 42.5 | 49.1 |
| Sustained disk to buffer | 36.6 | 42.3 |
| Buffer - Host transfer<br>68/80 pin FAST-40<br>68/80 pin Ultra 160/m | <br>80<br>160 | <br>80<br>160 |

Figure 12. Data transfer speed

The sustained customer data rate (in MB per second) takes into consideration a number of parameters, including head and cylinder switch times, to provide the customer with an average data transfer rate.

Calculate the data rates of 52.8 MB/s in zone 0 and 36.6 MB/s in zone 10 (based on the 36-GB model format) using the following formula:

$$\frac{\text{(Number of sectors per revolution) * (Bytes per sector)}}{\text{(Time per revolution) + (Head switch time)}}$$

**Note:** The number of sectors per track varies due to the usage of linear recording density.

The sustained disk to buffer transfer rate (in MB/s) takes into consideration a number of parameters, including head and cylinder switch times, to provide the customer with an average data transfer rate.

The maximum data transfer rate of the SCSI bus is the difference between the instantaneous buffer and the host transfer rate (in MB/s). This specification does not take the speed of the host into consideration.

## 4.3.10 Buffering operation (Read Ahead/Write Cache)

This hard drive has a data buffer for Read Ahead and Write Caching. The SCSI interface specification has details on the buffer.

## 4.3.11 Throughput

### 4.3.11.1 Simple sequential access

| Operation | Typical (ms) | Max (ms) |
|---|---|---|
| Sequential Read: Zone 0<br>at 465 sectors/ track | 83.4 | 85.0 |
| Sequential Read: ID zones<br>36-GB model at 322 sectors/track | 120 | 123 |
| Sequential Read: ID zones<br>18-GB model at 372 sectors/track | 104 | 106 |

Figure 13. Simple sequential access

The above table gives the time required to Read or Write a total of 8,000 consecutive sectors (4,096,000 Bytes) accessed by 128 Read/Write commands.

### 4.3.11.2 Random access

| **Operation: 36-GB Model** | **Typical** (sec) | **Max** (sec) |
|---|---|---|
| Random Read (Full Volume) | 3.4 | 3.6 |
| Random Write (Cache Disabled) | 3.9 | 4.1 |
| Random Write (Cache Enabled) | 3.3 | 3.5 |

| **Operation: 18-GB Model** | **Typical** (sec) | **Max** (sec) |
|---|---|---|
| Random Read (Full Volume) | 3.2 | 3.4 |
| Random Write (Cache Disabled) | 3.6 | 3.8 |
| Random Write (Cache Enabled) | 3.2 | 3.4 |

Figure 14. Random access

The above table gives the time required to execute a total of 1,000 Read/Write commands of 1K block transfers at 512 bytes per sector using a queue depth of 16.

## 4.3.12 Drive organization

### 4.3.12.1 Drive format

Media defects are remapped to the next available sector during the Format Process in manufacturing. An internally maintained table calculates mapping from the Logical Block Address (LBA) to the physical locations. Section 5.0Physical format" on page 17 describes the defect flagging process, a shipment prerequisite to ensure that the user will derive the maximum performance available from this drive design.

### 4.3.12.2 Cylinder allocation

**36-GB model**

| Data Zone | Physical Cylinder Number | Blocks / Track |
|---|---|---|
| 0 | 0 - 3276 | 465 |
| 1 | 3277 - 4730 | 454 |
| 2 | 4731 - 5590 | 442 |
| 3 | 5591 - 6728 | 434 |
| 4 | 6729 - 8331 | 413 |
| 5 | 8332 - 9036 | 403 |
| 6 | 9037 - 10205 | 387 |
| 7 | 10206 - 11957 | 372 |
| 8 | 11958 - 12768 | 351 |
| 9 | 12769 - 13742 | 336 |
| 10 | 13743 - 14532 | 322 |

**18-GB model**

| Data Zone | Physical Cylinder Number | Blocks / Track |
|---|---|---|
| 0 | 0 - 3276 | 465 |
| 1 | 3277 - 4730 | 454 |
| 2 | 4731 - 5590 | 442 |
| 3 | 5591- 6728 | 434 |
| 4 | 6729 - 8331 | 413 |
| 5 | 8332 - 9036 | 403 |
| 6 | 9037 - 10205 | 387 |
| 7 | 10206 - 10311 | 372 |
| 8 | NA | NA |
| 9 | NA | NA |
| 10 | NA | NA |

Figure 15. Cylinder allocation

**Note:** Calculations for physical cylinders begin from the starting data track, 0. Physical cylinders and logical cylinders do not necessarily correlate to one another.

**Data cylinder** These contain data which the user has saved and which Read or Write commands can access.

**Spare cylinder** Allocated spare cylinders accommodate user data reassigned from any defect location on the disk. Spare cylinders are distributed every 256 cylinders, starting from physical cylinder number 0, to minimize Seek distance to a reassigned sector.

This page intentionally left blank.

# 5.0 Physical format

Media defects are remapped to the next available sector during the Format Process in manufacturing. All manufacturing located defects are itemized and mapped on the drive in the Primary Defect List (PList). The defects are then eliminated from the user data area.

Defects identified after manufacturing are flagged and logged as Grown Defect List (GList) entries. An internal table reassigns Logical Block Addresses (LBAs) to accommodate the relocation of data from the affected areas.

## 5.1 Shipped format (PList)

- Reassigned sectors occupy spare tracks distributed every 256 cylinders.
- No data sectors are used for sparing purposes.

| N | N+1 | Defect | N+2 | Defect | N+3 |
|---|---|---|---|---|---|
| | | (Skip) | | (Skip) | |

Figure 16. Shipped format (PList)

**Note:** Defects are skipped without any constraint, such as track or cylinder boundary. An internal table automatically calculates from LBA to physical.

## 5.2 Reassigned format (GList)

- GList is prepared for 3,279 LBAs
- Re-re-assign of the same LBA does not increase GList entry
- A cylinder for spare sectors is prepared every 256 physical cylinders

**Note:** Normal hard disk drive maintenance includes GList entries. GList entries are possible during early drive usage and are primarily caused by handling.

This page intentionally left blank.

# 6.0 Specification

## 6.1 Electrical interface specification

### 6.1.1 Power connector

The DC power connectors (68- and 80-pin SCA-2) are an integral portion of the 68-pin SCSI ‘Unitized’ Connectors and 80-pin 'Single Connector Attachment' (SCA-2) connectors.

The 8/16 SCSI bus offered by the 68-pin connector models is compatible with the ANSI SCSI “P” connector, which supports wide data transfers. Figure 17 below shows the pin assignment for the 68-pin power connector.

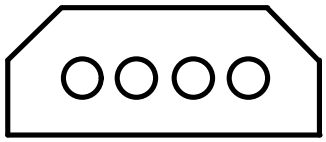

+5V GND GND +12V

Mating side view

Figure 17. Power connector pin assignment

The 80-pin SCA-2 models use a DDK connector (PN HD2-PA080-A14B) or equivalent version compatible with the specifications of "Single Attachment for Small SCSI Disk Drives," SPI-3 document, Annex C.

The user must ensure that the electrical and mechanical characteristics of the mating connector are fully compatible with the disk drive.

### 6.1.2 SCSI bus connector

68-pin and 80-pin SCSI connectors support the following:

| Interface | 68-pin | 80-pin (SCA-2) |
|---|---|---|
| LVD / SE Multi-mode | Yes | Yes |

Figure 18. SCSI bus connector

### 6.1.2.1 68-pin LVD/SE multi-mode

| Pin | Signal | | Pin | Signal |
|---|---|---|---|---|
| 01 | +DB(12) | | 35 | -DB(12) |
| 02 | +DB(13) | | 36 | -DB(13) |
| 03 | +DB(14) | | 37 | -DB(14) |
| 04 | +DB(15) | | 38 | -DB(15) |
| 05 | +DB(P1) | | 39 | -DB(P1) |
| 06 | +DB(0) | | 40 | -DB(0) |
| 07 | +DB(1) | | 41 | -DB(1) |
| 08 | +DB(2) | | 42 | -DB(2) |
| 09 | +DB(3) | | 43 | -DB(3) |
| 10 | +DB(4) | | 44 | -DB(4) |
| 11 | +DB(5) | | 45 | -DB(5) |
| 12 | +DB(6) | | 46 | -DB(6) |
| 13 | +DB(7) | | 47 | -DB(7) |
| 14 | +P_CRCA | | 48 | -P_CRCA |
| 15 | Ground | | 49 | Ground |
| 16 | DIFFSENS (*2) | | 50 | Ground |
| 17 | TERMPWR (*1) | | 51 | TERMPWR (*1) |
| 18 | TERMPWR (*1) | | 52 | TERMPWR (*1) |
| 19 | RESERVED | | 53 | RESERVED |
| 20 | Ground | | 54 | Ground |
| 21 | +ATN | | 55 | -ATN |
| 22 | Ground | | 56 | Ground |
| 23 | +BSY | | 57 | -BSY |
| 24 | +ACK | | 58 | -ACK |
| 25 | +RST | | 59 | -RST |
| 26 | +MSG | | 60 | -MSG |
| 27 | +SEL | | 61 | -SEL |
| 28 | +C/D | | 62 | -C/D |
| 29 | +REQ | | 63 | -REQ |
| 30 | +I/O | | 64 | -I/O |
| 31 | +DB(8) | | 65 | -DB(8) |
| 32 | +DB(9) | | 66 | -DB(9) |
| 33 | +DB(10) | | 67 | -DB(10) |
| 34 | +DB(11) | | 68 | -DB(11) |

Figure 19. Table of signals for 68-pin connector

**Note:** *1 TERMPWR can be disabled.
*2 HVD is not supported.

The 8-bit devices which connect to the 68-pin connector should have the following signals inactive (high):

-DB(8), -DB(9), -DB(10), -DB(11), -DB(12), -DB(13), -DB(14), -DB(15), -DB(P1)

Connect all other signals as defined.

### 6.1.2.2 80-pin (SCA-2) LVD/SE multi-mode

| Pin | Signal | | Pin | Signal |
|---|---|---|---|---|
| 01 | 12 V charge | | 41 | 12 V ground |
| 02 | 12 V | | 42 | 12 V ground |
| 03 | 12V | | 43 | 12 V ground |
| 04 | 12 V | | 44 | Mated 1 |
| 05 | NC | | 45 | NC |
| 06 | NC | | 46 | DIFFSNS (*1) |
| 07 | -DB(11) | | 47 | +DB(11) |
| 08 | -DB(10) | | 48 | +DB(10) |
| 09 | -DB(9) | | 49 | +DB(9) |
| 10 | -DB(8) | | 50 | +DB(8) |
| 11 | -I/O | | 51 | +I/O |
| 12 | -REQ | | 52 | +REQ |
| 13 | -C/D | | 53 | +C/D |
| 14 | -SEL | | 54 | +SEL |
| 15 | -MSG | | 55 | +MSG |
| 16 | -RST | | 56 | +RST |
| 17 | -ACK | | 57 | +ACK |
| 18 | -BSY | | 58 | +BSY |
| 19 | -ATN | | 59 | +ATN |
| 20 | -P_CRCA | | 60 | +P_CRCA |
| 21 | -DB(7) | | 61 | +DB(7) |
| 22 | -DB(6) | | 62 | +DB(6) |
| 23 | -DB(5) | | 63 | +DB(5) |
| 24 | -DB(4) | | 64 | +DB(4) |
| 25 | -DB(3) | | 65 | +DB(3) |
| 26 | -DB(2) | | 66 | +DB(2) |
| 27 | -DB(1) | | 67 | +DB(1) |
| 28 | -DB(0) | | 68 | -DB(0) |
| 29 | -DB(P1) | | 69 | +DB(P1) |
| 30 | -DB(15) | | 70 | +DB(15) |
| 31 | -DB(14) | | 71 | +DB(14) |
| 32 | -DB(13) | | 72 | +DB(13) |
| 33 | -DB(12) | | 73 | +DB(12) |
| 34 | 5 V | | 74 | Mated 2/GND |
| 35 | 5 V | | 75 | 5 V ground |
| 36 | 5V charge | | 76 | 5 V ground |
| 37 | NC | | 77 | Active LED out |
| 38 | RMT START | | 78 | DLYE_START |
| 39 | SCSI ID(0) | | 79 | SCSI ID(1) |
| 40 | SCSI ID(2) | | 80 | SCSI ID(3) |

Figure 20. Table of signals for 80-pin connector

**Note:** *1 HVD is not supported.

**Note:** The SCA-2 connector is intended for direct back plane attachment and not for cable attachment to the bus. The SCA-2 connector is not mechanically compatible with either the 50-pin "A" or the 68-pin "P" connectors, as defined in the ANSI SCSI standard.

The 8-bit devices which connect to the SCA-2 connector should have the following signals inactive (high):

–DB(8), – DB(9), – DB(10), –DB(11), –DB(12), –DB(13), –DB(14), –DB(15), –DB(P1)

Connect all other signals as defined.

## 6.1.3 SCSI bus cable and drive configurations

The drive operates in all of the STA-endorsed SCSI bus configurations shown in the figure below.

| STA Term For Interface Type | Bus Speed [MB/s (Max.)] | Bus Width in bits | Max. Bus Length in Meters[1] | | | Number of Devices Supported (Max.) |
|---|---|---|---|---|---|---|
| | | | SE | LVD | HVD[2] | |
| **SCSI-1** | 5 | 8 | 6 | ND+ | NA | 8 |
| **Fast SCSI** | 10 | 8 | 3 | ND+ | NA | 8 |
| **Fast Wide SCSI** | 20 | 16 | 3 | ND+ | NA | 16 |
| **Ultra SCSI** | 20 | 8 | 1.5 | ND+ | NA | 8 |
| **Ultra SCSI** | 20 | 8 | 3 | - | - | 4 |
| **Wide Ultra SCSI** | 40 | 16 | - | ND+ | NA | 16 |
| **Wide Ultra SCSI** | 40 | 16 | 1.5 | - | - | 8 |
| **Wide Ultra SCSI** | 40 | 16 | 3 | - | - | 4 |
| **Ultra2 SCSI** | 40 | 8 | NA | 12 | NA | 8 |
| **Wide Ultra2 SCSI** | 80 | 16 | NA | 12 | NA | 16 |
| **Ultra 160 SCSI** | 160 | 16 | NA | 12 | ND | 16 |

Figure 21. SCSI bus cable and drive configurations

**Notes:** [1] The maximum bus lengths listed may be exceeded in point-to-point applications.
[2] The Ultrastar 36Z15 drives do not support High Voltage Differential (HVD).

**Abbreviations:**

NA SE is not available for speeds beyond Ultra.
ND The original SCSI standard did not define the specification for the item.
ND+ ND definition plus, if all devices on the bus support Low Voltage Differential (LVD), then operation at a 12-meter bus length is possible at the indicated bus speed. However, if any SE-only device is attached to the bus, the bus defaults to SE mode, and SE specifications apply.

## 6.1.4 SCSI bus terminators

No active SCSI termination exists on the drive. The using system is responsible for ensuring termination of all required signals at both the beginning (i.e., host adapter) and end of the bus (i.e., last device). Termination must be present within four inches of the physical ends of the SCSI bus and external to the drive. For maximum performance and reliability, implement LVD termination in the system configuration.

## 6.1.5 SCSI bus termination power

The 5V power supply on the 68-pin drive models may provide termination power, or “term power,” as an option through a current limiter and Shottky diode. In order to use termination power, the user must install a jumper between pins 3 and 4 of the Option Jumper Block (J4). Install this jumper only on the device which is physically closest to the terminator.

80-pin SCA-2 models do not support SCSI bus termination power.

## 6.1.6 SCSI bus electrical characteristics

The ANSI document, SCSI Parallel Interface-3 T10/1302 revision 11, defines all electrical characteristics of the bus.

**Note:** If the drive is connected to an HVD bus, the drive I/O will sustain permanent damage.

## 6.1.7 Hot plug/unplug

The term ‘Hot Plug’ refers to the action of mechanically engaging a device to the power or bus or both while other devices on the same bus may be active. A comprehensive description on the state of the SCSI bus during this event is available in the SCSI-3 Parallel Interface Standard.

While every effort was made to design the drive to not influence the SCSI bus during these events, it is the responsibility of the system to regulate voltage and to conform to specified operational and non-operational shock limits of the drive being handled, as well as those drives adjacent to it. If there is concern that the operational shock limits cannot be maintained for the drive involved during a Hot Unplug operation, issue a SCSI Stop Unit command to the drive and allow it to complete the command. It is also recommended that the subject drive rest in the drive bay for a minimum of 15 seconds after completing the Stop Unit command before being unplugged. This is important since the operational shock limits are in effect whenever the drive is operational or in a spin down state.

The power supply ripple on adjacent operational drives should not exceed the ±5% regulation tolerance during either Hot Plug or Unplug events.

## 6.1.8 The 68-pin auxiliary connector

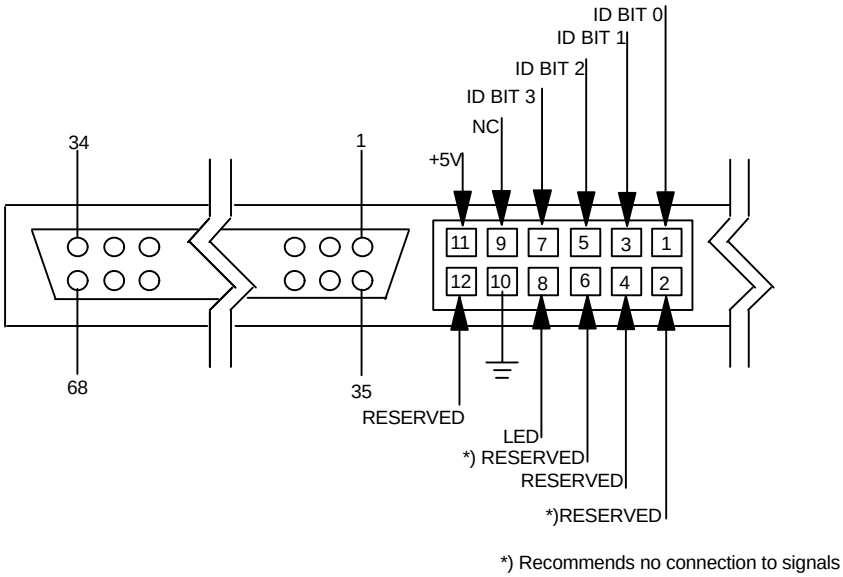


Figure 22. The 68-pin auxiliary connector

# 6.2 Option jumper block

Two option blocks designated as J4 and J6 are located on the card of 68- and 80-pin models as shown in Figure 23 below.

The J4 option block has 14 jumper positions, including drive ID bit settings and the control for Terminator Power supply.

In 68-pin models pin 3 of J4 is connected to the 5V line via Polyswitch. In 80-pin models the pin is not connected.

The J6 option block also has 14 jumper positions which are primarily for motor and unit control options.

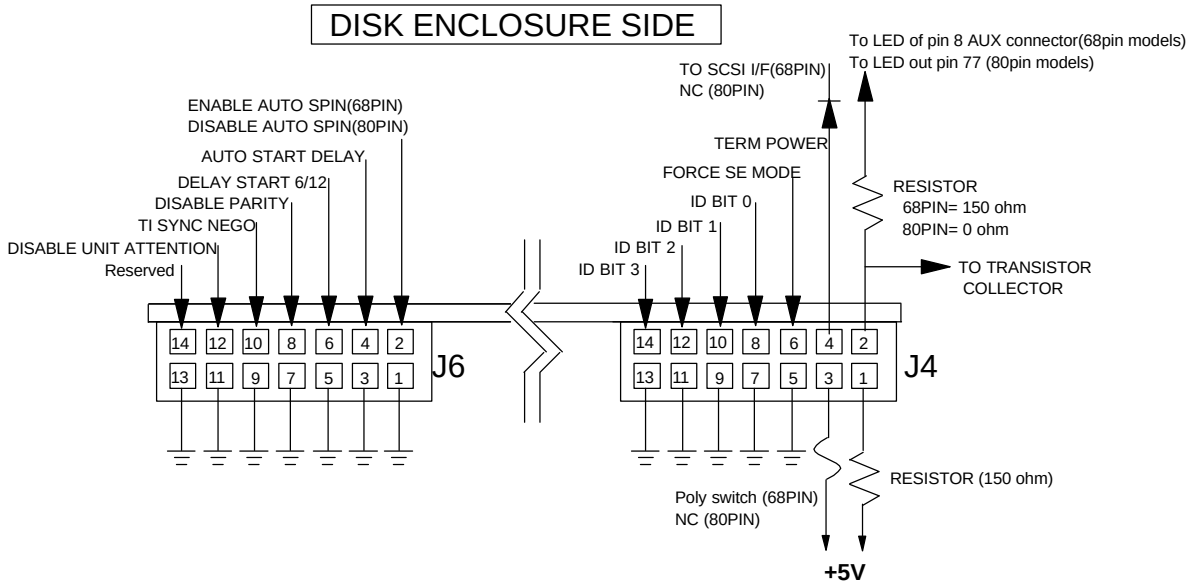


Figure 23. Option jumper block

## 6.2.1 Jumper signal description on option block J6

### 6.2.1.1 Position #1-2 - Enable Auto Spin Up - 68-pin model

If a jumper is installed, the drive will spin up automatically after power on reset. If a jumper is not installed, the drive will not spin up unless a START UNIT command is received.

### 6.2.1.2 Position #1-2 - Disable Auto Spin Up - 80-pin model

If a jumper is not installed, the drive will spin up automatically after power on reset. If a jumper is installed, the drive will not spin up unless a START UNIT command is received.

### 6.2.1.3 Position #3-4 and #5-6 - Auto Start Delay & Delay Start 6/12

These pins control when and how the drive spins up in conjunction with Position #1-2 on J6. See Figure 24 below. When both Auto Spin up and Auto Start Delay are enabled, the start of the drive will be delayed by a period of time multiplied by its own SCSI address. If Auto Spin up is disabled, these jumpers

will be ignored. Placing a jumper on delay start 6/12 results in a start up delay of 12 seconds times the SCSI ID.

*Note: In the table below 'on' means a shunt jumper is installed and 'off' means that a shunt jumper is not installed.*

**68-pin models**

| Auto Start Delay | Auto Spin Up | Delay Start 6/12 | Auto Start | Delay Multiplier |
|---|---|---|---|---|
| off | off | off | NO | - |
| off | on | off | YES | 0 |
| on | off | off | NO | - |
| on | on | off | YES | 6 |
| off | off | on | NO | - |
| off | on | on | YES | 0 |
| on | off | on | NO | - |
| on | on | on | YES | 12 |

**80-pin models**

| Auto Start Delay | Auto Spin Up | Delay Start 6/12 | Auto Start | Delay Multiplier |
|---|---|---|---|---|
| off | off | off | YES | 0 |
| off | on | off | NO | - |
| on | off | off | YES | 6 |
| on | on | off | NO | - |
| off | off | on | YES | 0 |
| off | on | on | NO | - |
| on | off | on | YES | 12 |
| on | on | on | NO | - |

Figure 24. Auto Start Delay & Delay Start 6/12 drive behavior

### 6.2.1.4 Position #7-8 - Disable SCSI Parity Check

Installing a jumper shunt disables SCSI Parity checking.

### 6.2.1.5 Position #9-10 - Enable TI-SDTR

Installing a shunt jumper enables Target Initiated Synchronous Data Transfer Request Negotiation. For additional negotiation controls refer to the paragraph describing the SSM bit in Section 7.7.2 on page 87.

### 6.2.1.6 Position #11-12 on J6 - Disable Unit Attention

Installing a shunt jumper enables control of UAI (Unit Attention Inhibit) bit in Mode Page 0.

### 6.2.1.7 Position #13-14 on J6

Reserved.

## 6.2.2 Jumper signal description on option block J-4

### 6.2.2.1 Position #1-2 - LED pins

The LED pins are used to drive an external Light Emitting Diode. Up to 30 mA of sink current capability is provided. The LED Anode must be tied to the current limited + 5 V source provided on the pin for LED Anode at the Location #1 on J4 jumper block. The LED Cathode is then connected to the pin for LED Cathode at the Location #2 on J4 jumper block to complete the circuit. See Figure 25 below.

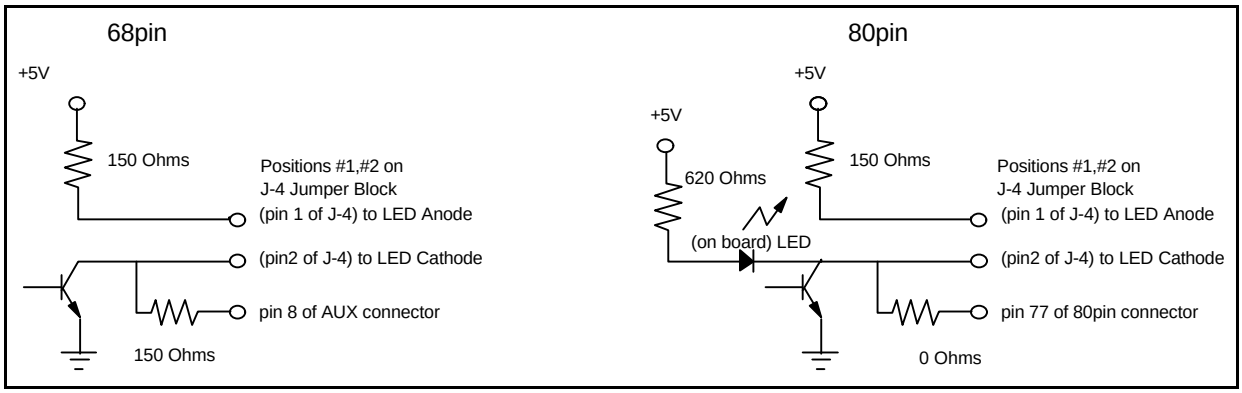


Figure 25. LED circuit diagram

### 6.2.2.2 Position #3-4 - Term Power

Not connected for 80-pin models. If a jumper is installed on 68-pin models, termination power is supplied to pins 17, 18, 51, and 52 of the 68 pin of the SCSI interface.

### 6.2.2.3 Position #5-6 - Force SE mode

If a jumper is installed, the drive functions as a single-ended mode drive.

### 6.2.2.4 Position #7-8 to #13-14 - SCSI ID

These four lines (–DAS0, –DAS1, –DAS2, –DAS3) define the drive's ID on the SCSI BUS. –DAS0 is the least significant bit and –DAS3 is the most significant bit. Device ID is defined as shown in the table below. 'On' means a jumper is installed and 'off' means that no jumper is installed.

| -DAS3 (1) | -DAS2 (2) | -DAS1 (3) | -DAS0 (4) | Device ID |
|---|---|---|---|---|
| off | off | off | off | 0 - shipping default of 80 pin |
| off | off | off | on | 1 |
| off | off | on | off | 2 |
| off | off | on | on | 3 |
| off | on | off | off | 4 |
| off | on | off | on | 5 |
| off | on | on | off | 6 - shipping default of 68 pin |
| off | on | on | on | 7 |
| on | off | off | off | 8 |
| on | off | off | on | 9 |
| on | off | on | off | 10 |
| on | off | on | on | 11 |
| on | on | off | off | 12 |
| on | on | off | on | 13 |
| on | on | on | off | 14 |
| on | on | on | on | 15 |

Figure 26. SCSI device ID

# 6.3 Environmental requirements

## 6.3.1 Temperature and humidity, operating

| | |
|---|---|
| Temperature | 5 to 55°C (see note below) |
| Relative humidity | 8 to 90%, non-condensing |
| Maximum wet bulb temperature | 29.4°C, non-condensing |
| Maximum temperature gradient | 15°C/Hour |
| Altitude | –300 to 3,048 meters |

Figure 27. Operating temperature and humidity

The maximum allowable HDA case temperature is 65°C, as measured at the locations defined below. Operation of the drive at the maximum case temperature is intended only for short periods of time. Continuous operation at elevated temperatures will reduce product reliability. Maintain non- condensing conditions at all times.

Dimensions are in mm.

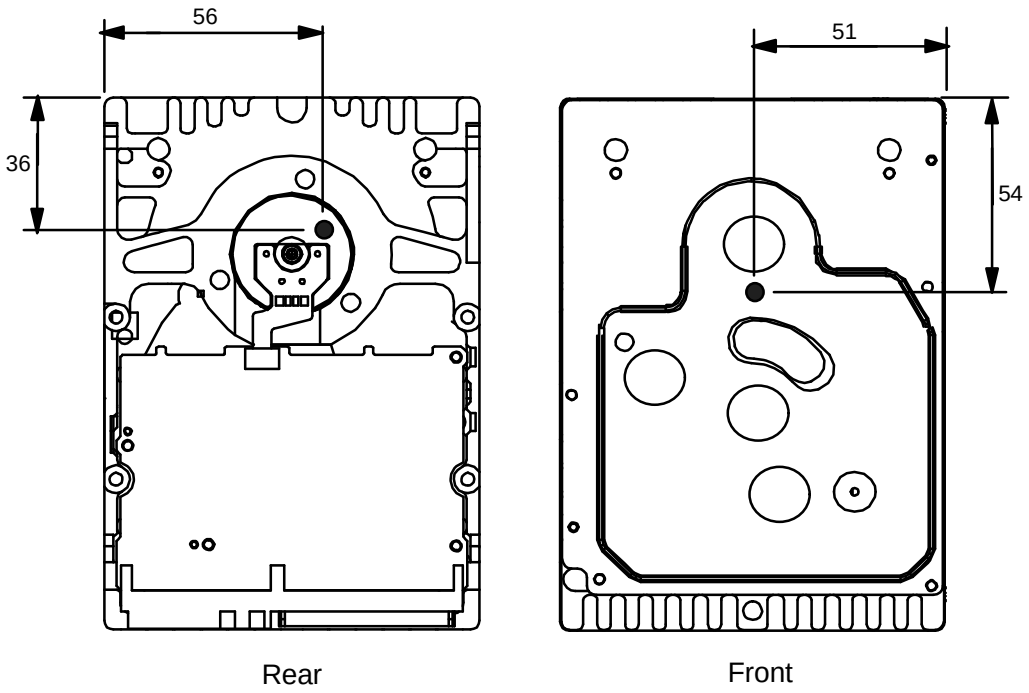


Figure 28. Drive Temperature Measurement Points

## 6.3.2 Temperature and humidity, shipping

| | |
|---|---|
| Temperature | –40 to 70°C |
| Relative humidity | 5 to 95%, non-condensing |
| Maximum wet bulb temperature | 35°C, non-condensing |
| Maximum temperature gradient | 15°C/Hour |
| Altitude | –300 to 12,000 meters |

Figure 29. Shipping temperature and humidity

## 6.3.3 Temperature and humidity, storage

| | |
|---|---|
| Temperature | 0 to 65°C |
| Relative humidity | 5 to 95%, non-condensing |
| Maximum wet bulb temperature | 35°C, non-condensing |
| Altitude | –300 to 12,000 meters |

Figure 30. Storage temperature and humidity

**Note:** The maximum storage period in the shipping package is one year.

## 6.3.4 Electrical component cooling

Drive component temperatures must remain within the limits specified in Figure 31 and must not be exceeded under any operating condition. The maximum allowable HDA case temperature to meet the reliability expectation for the Ultrastar 36Z15 drive is 55°C, which may require the drive to have forced air cooling.

| **Component Module** | **Location** | **Maximum allowable surface temperature** (degrees C) | **Recommended operating surface temperature** (degrees C) |
|---|---|---|---|
| MPU + HDC integration module | 1 | 95 | 75 |
| DRAM | 2 | 90 | 70 |
| VCM & spindle driver | 3 | 95 | 75 |
| Channel module | 4 | 95 | 75 |

Figure 31. Electrical component cooling

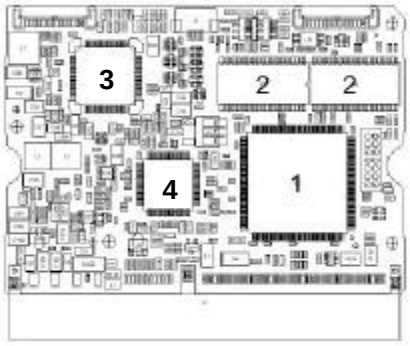


Figure 32. Module location

## 6.3.5 Corrosion susceptibility

The hard disk drive will not show any sign of corrosion inside or outside the HDA and remains functional to specification when tested at 50°C temperature and 90% relative humidity for a period of seven days.

# 6.4 DC power requirements

The following voltage specifications apply at the drive power connector. For the SCA 80-pin model, hot plug functions are supported. Make connections to the drive in a safety extra low voltage (SELV) circuit. No special power on/off sequencing is required.

Adequate secondary over-current protection is the responsibility of the system integrator. For safety purposes, a 10-A limit is required.

## 6.4.1 Input voltage

| | **During normal operation and spin up** | **Absolute max spike voltage** (to avoid damage to drive electronics) |
|---|---|---|
| +5 volt supply | ±5% | 7 V |
| +12 volt supply | ±5%* | 15 V |

Figure 33. Input voltage

*A deviation up to –8% is acceptable during drive spin up, however, spin up time is not guaranteed.

## 6.4.2 Power dissipation and peak current

All power numbers specified represent typical values under nominal temperature and voltage conditions.

| **36-GB Models** | **+5 Volts mean (Amps)** | **+12 Volts mean (Amps)** | **Total (Watts)** |
|---|---|---|---|
| Idle average | 0.5 | 0.9 | 13.5 |
| Seek average (*1) | 0.7 | 1.2 | 18.0 |
| Start up (max) | 1.0 | 2.5 | N/A |
| Random R/W peak (*2) | 1.3 | 2.7 | N/A |
| Random R/W average (*2) | 0.6 | 1.2 | 17.0 |

| **18-GB Models** | **+5 Volts mean (Amps)** | **+12 Volts mean (Amps)** | **Total (Watts)** |
|---|---|---|---|
| Idle average | 0.5 | 0.7 | 11.0 |
| Seek average (*1) | 0.7 | 1.0 | 16.0 |
| Start up (max) | 1.0 | 2.5 | N/A |
| Random R/W peak (*2) | 1.3 | 2.5 | N/A |
| Random R/W average (*2) | 0.6 | 1.0 | 15.0 |

Figure 34. Power dissipation and peak current

**Duty Cycle:** *1 Seek duty 100% with no delay between Seeks
*2 Random Write/Read: 120 ops, 4 Kbytes

During drive startup and Seek, the drive generates 12-volt ripple (referred to as dynamic loading). If the power of several drives is daisy-chained together, the power supply ripple plus the dynamic loading of the drive must remain within the regulation tolerance of ±5%. A common supply which supports separate power leads to each drive is a more desirable method of power distribution.

To prevent external electrical noise from interfering with drive performance, the drive must be mounted securely by four screws to the user system frame. There must be no electrical level difference at any of the four screw positions in the system frame. The peak-to-peak level difference between any of the four screw positions and the drive power connector ground must be less than ±300 millivolts.

## 6.4.3 Power current profiles

The following are sample 12V and 5V current profiles. On all figures, the top trace is 12V current and the bottom trace is 5V current.

### 6.4.3.1 Start up current

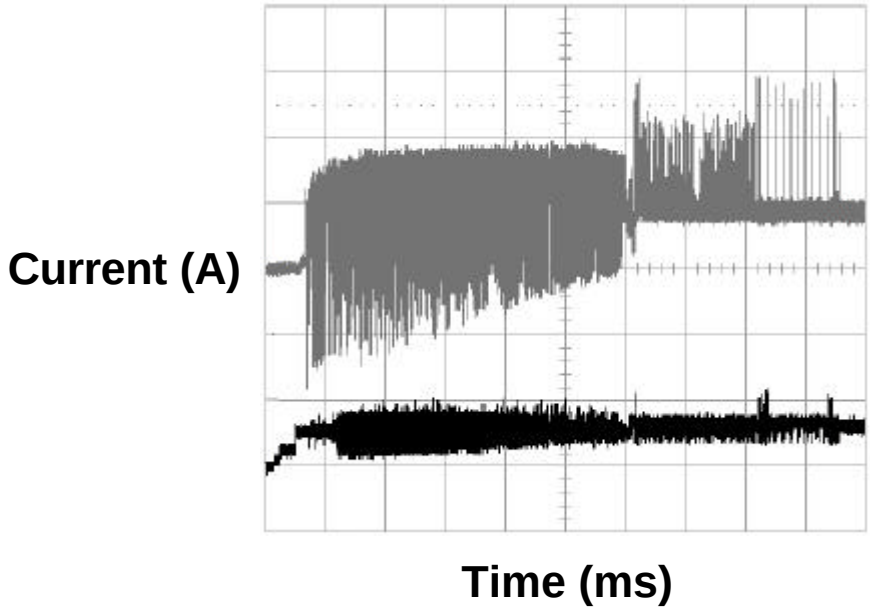


Figure 35. Start up current

### 6.4.3.2 Idle current

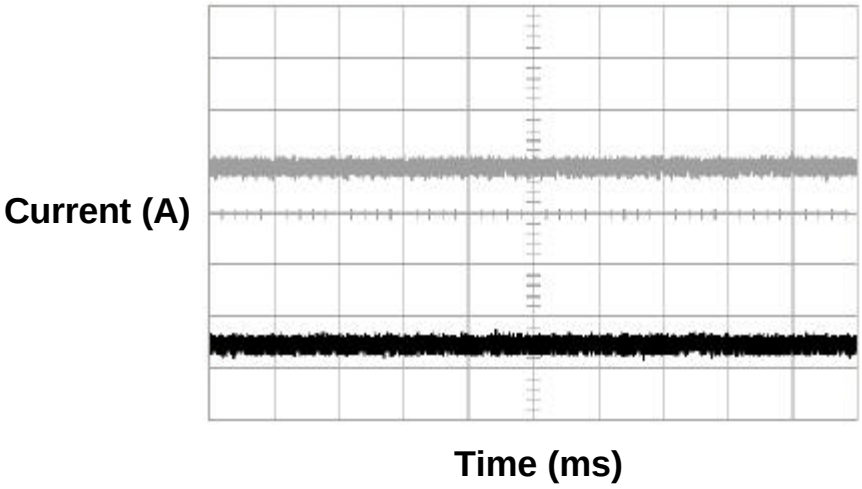


Figure 36. Idle current

### 6.4.3.3 Random read/write current

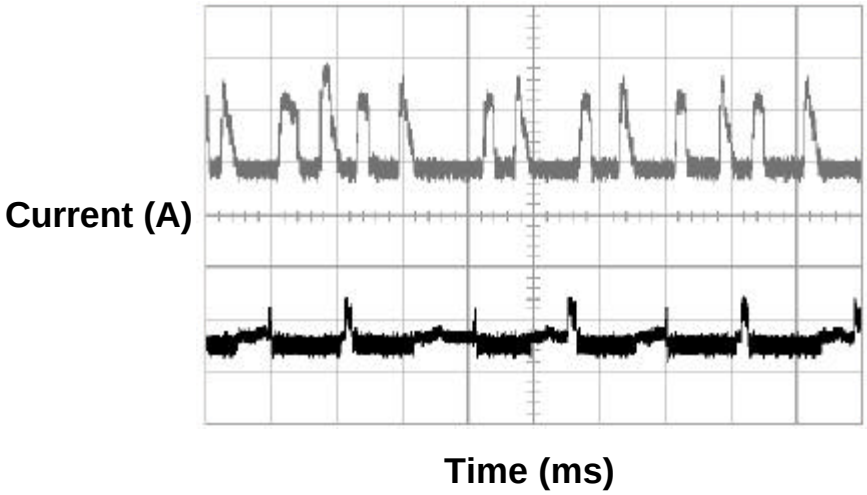


Figure 37. Random read/write current

### 6.4.3.4 Power consumption efficiency index

Power consumption efficiency index values are not required for the Ultrastar 36Z15 since it is an enterprise class drive. It has been designed to perform within the efficiency index values applicable to 15,000-RPM hard disk drives if required.

# 6.5 Reliability

## 6.5.1 Data integrity

No more than one sector is lost at power loss condition during Write operation when the Write Cache option is disabled. If the Write Cache option is active, the data in Write Cache will be lost. Verify completion of the last Write access before powering off.

## 6.5.2 Start/stop cycles

The drive is designed to support 50,000 start/stop cycles at 55°C HDA case temperature and a minimum of 10,000 start/stop cycles at extreme temperature as specified in section 6.3, “Environmental requirements,” on page 28.

## 6.5.3 Drive life

The life of the drive is five years under the following conditions:

- 43,800 POH (730 POH per month)
- Seeking/Writing/Reading operations to constitute 20% of POH at 55°C HDA case temperature
- Maximum HDA case temperature of 65°C

The information provided is not a guarantee of actual drive performance but is the best estimate based on typical work loads, operating conditions, and system environment.

## 6.5.4 Preventive maintenance

None required.

## 6.5.5 Data reliability

The probability of uncorrectable data error rate is 10 in $10^{16}$ bits when measured at nominal DC voltage conditions and recommended operating temperatures with Full ERP enabled. See Section 6.3.1, "Operating temperature and humidity," on page 28.

The following ECC on the fly correction is implemented:

- 1 symbol: 8 bits
- 3 interleave
- 15 symbols, 5 per interleave for on the fly correction in burst or random modes. Software correction up to 30 bytes with pointers

# 6.6 Mechanical specifications

## 6.6.1 Physical dimensions

| | |
|---|---|
| Height | 25.7 ± 0.4mm |
| Width | 101.6 ± 0.4mm |
| Length | 146.0 ± 0.6mm |
| Weight (max)<br>36-GB models:<br>18-GB models: | <br>0.75 kilogram<br>0.74 kilogram |

Figure 38. Physical dimensions

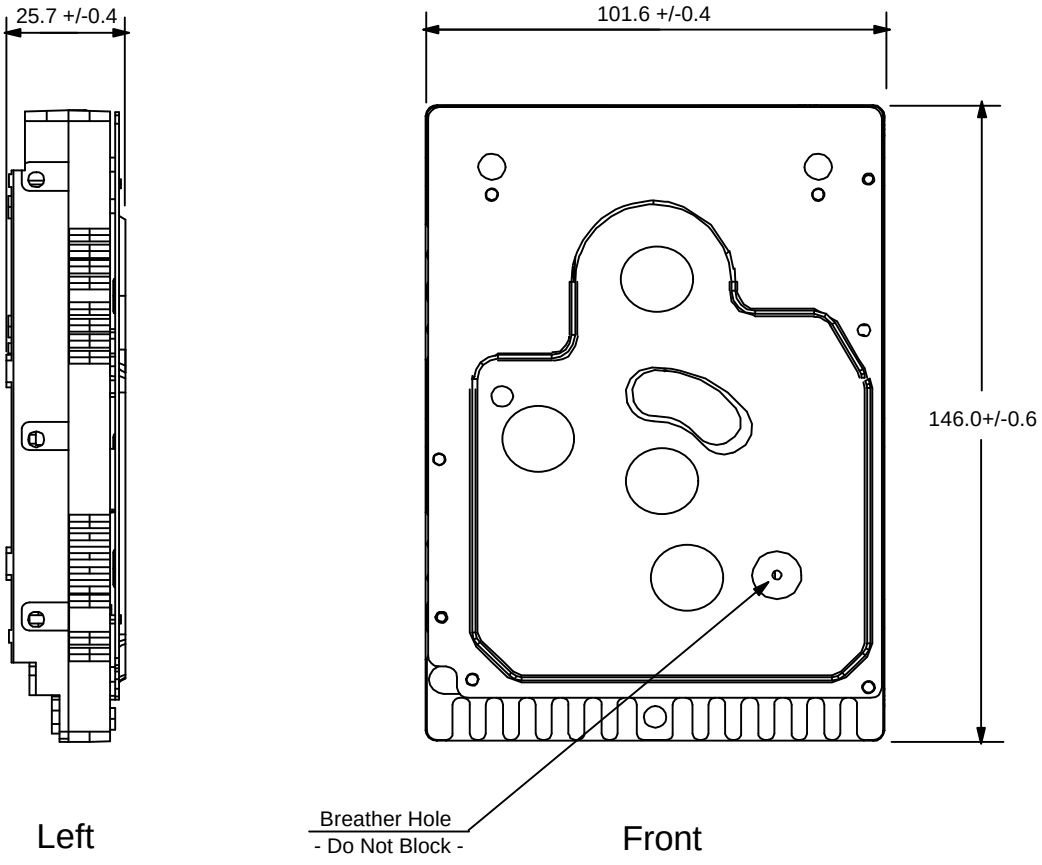


Figure 39. Physical dimensions illustrated

## 6.6.2 Drive mounting

The drive will operate in all axes (six directions). The drive may be sensitive to user mounting implementation due to frame distortion effects. To avoid performance degradation, mount the drive in the system securely enough to prevent excessive motion or vibration of the drive during Seek operations or spindle rotation, using appropriate screws or equivalent mounting hardware. IBM provides technical support to users who wish to further investigate details of their actual application.

Conduct drive level vibration and shock testing with the drive mounted to the test table by the four bottom screw holes.

All dimensions are in millimeters.

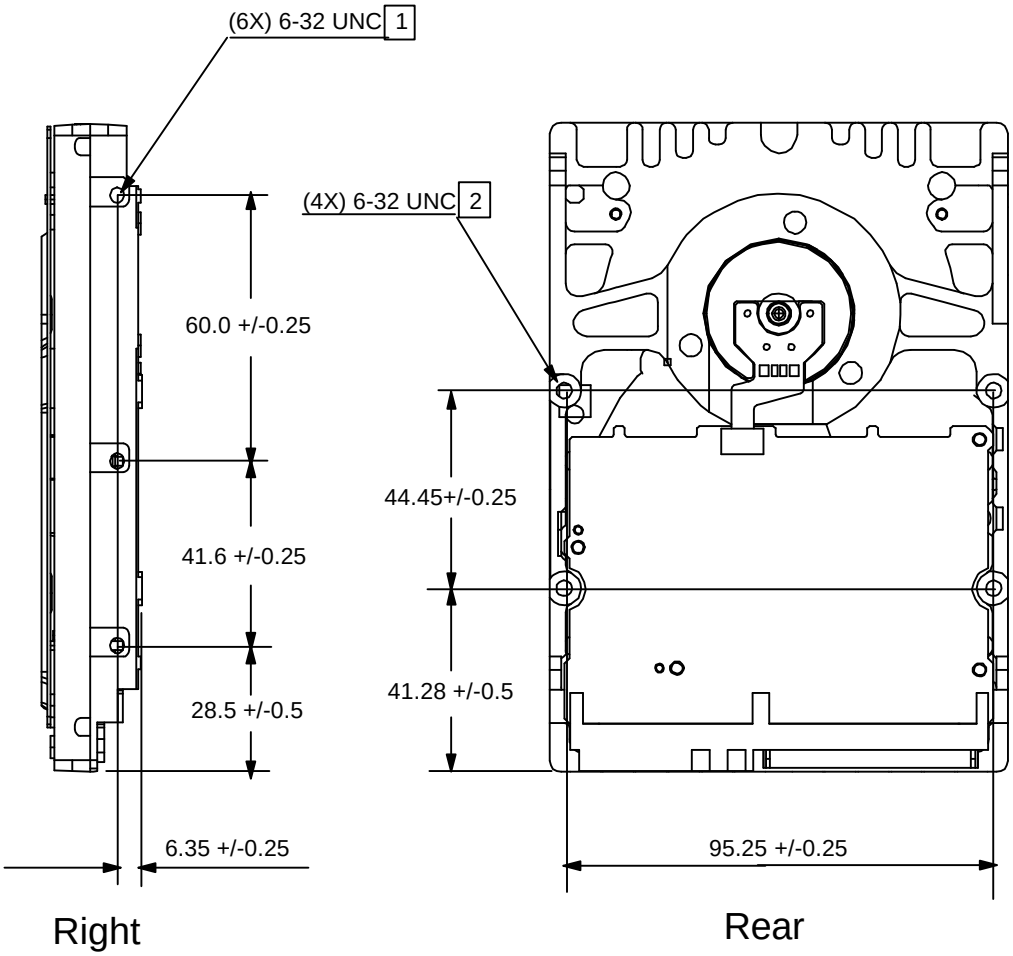


Figure 40. Mounting positions and tappings

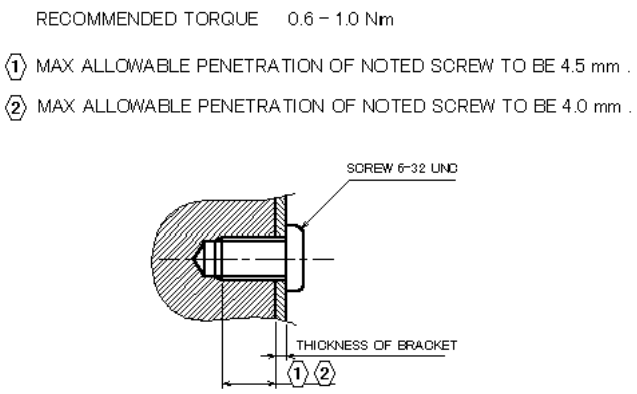


Figure 41. Maximum screw penetration

## 6.6.3 Dimensions of the 68-pin connector

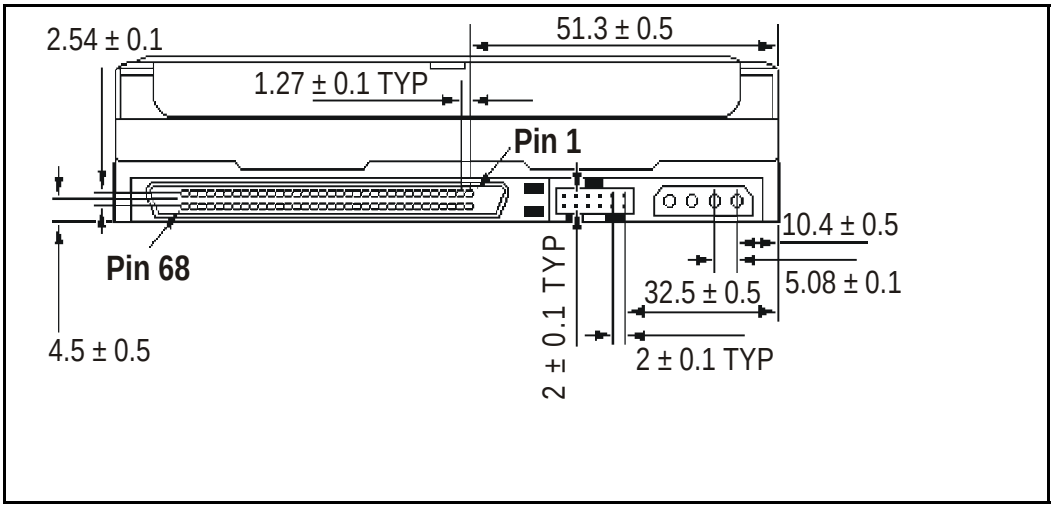


Figure 42. Dimensions of the 68-pin connector

## 6.6.4 Dimensions of the 80-pin connector

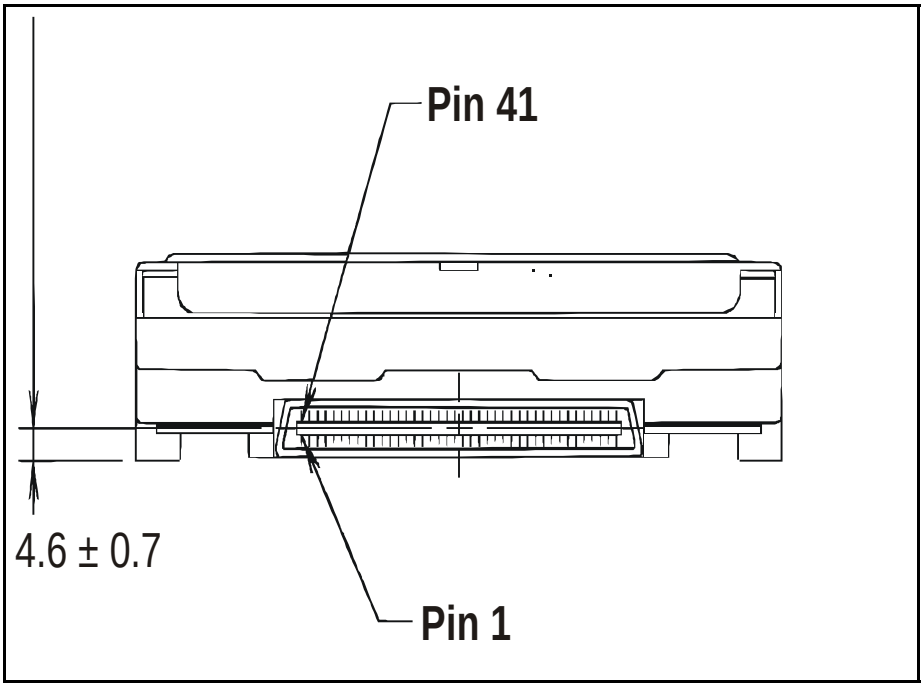


Figure 43. Dimensions of 80-pin connector

## 6.6.5 Heads unload and actuator lock

The drive moves heads out from the disks (unload) to the ramp to protect the disk data during shipping, moving, or storage. Upon power down, the drive automatically unloads the heads from the disk area. An inertia latch is engaged if the drive is subjected to a shock event. This causes the actuator to be securely locked while on the ramp.

# 6.7 Vibration

## 6.7.1 Random vibration, operating

The drive is designed to operate without permanent hardware damage and data loss when subjected to vibration levels not exceeding the vibration levels defined in Figure 39.

Measurements are taken at the mounting points of the drive during 30 minutes of random vibration using the power spectral density (PSD) levels as follows:

| Item | 5 Hz | 500 Hz | RMS |
|---|---|---|---|
| Random Vibration | 4.6 E-3 | 4.6 E-3 | 1.50 |
| Units | $\frac{G^2}{Hz}$ | $\frac{G^2}{Hz}$ | Grms |

Figure 44. Operating random vibration

## 6.7.2 Swept sine vibration, operating

The drive will operate without permanent hardware damage and data loss when subjected to sine vibration as follows:

No errors under the following conditions:

- 0.5 G 0-peak, 5–500–5 Hz sine wave
- 0.5 octave/minute sweep rate
- 3-minute dwell at two major resonance points

No data loss under the following conditions:

- 1.0 G 0-peak, 5–500–5 Hz sine wave
- 0.5 octave/minute sweep rate
- 3-minute dwell at two major resonance points

## 6.7.3 Random vibration, nonoperating

This test consists of applying random vibration in three mutually perpendicular axes at a time duration of ten minutes per axis. The PSD levels for the test simulates the shipping and relocation environment specified below. The drive does not sustain permanent hardware damage nor loss of previously recorded data after being subjected to the environment specified below.

| Item | 5 Hz | 500 Hz | RMS |
|---|---|---|---|
| Random Vibration | 1.0 E-2 | 1.0 E-2 | 2.23 |
| Units | $\frac{G^2}{Hz}$ | $\frac{G^2}{Hz}$ | Grms |

Figure 45. Nonoperating random vibration

## 6.7.4 Swept sine vibration, nonoperating

The drive will operate without permanent hardware damage and data loss when subjected to the following swept sine vibration environment:

- 2 G 0-peak, 5–500–5 Hz sine wave
- 0.5 octave/minute sweep rate
- 3 minute dwells at two major resonance

# 6.8 Shock

## 6.8.1 Operating shock

The disk drive meets IBM Standard C-S 1-9711-007 for the S5 product classification. The drive incurs no permanent hardware damage nor data loss when subjected to the following shock pulses:

- 15-G half sine wave shock pulse of 11-ms duration
- 60-G half sine wave shock pulse of 2-ms duration

The shock test consists of 10 shock pulses applied in each axis and direction for total of 60 pulses. The drive requires a delay between shock pulses to complete all necessary error recovery procedures (approximately a 30 second delay).

## 6.8.2 Square shock, nonoperating

The drive withstands the following square shock pulse without permanent hardware damage or data loss:

- Accelerating level 50 G
- Duration time 11 ms
- Rise and fall time 1 ms
- Minimum velocity change 4.57 meters/second

The shocks are applied in each direction of the drive for three mutually perpendicular axes and one axis at a time. Input levels are measured on a base plate to which the drive is attached with four screws. The shock inputs will not displace the heads from the actuator latch.

## 6.8.3 Half-sine shock, nonoperating

The drive incurs no permanent damage nor data loss when subjected to the following shock pulses:

- 80-G half sine wave shock pulse of 11-ms duration
- 250-G half sine wave shock pulse of 2-ms duration

All shock input is applied in the manner described in Section 6.8.2 above. The shock inputs will not displace the heads from the actuator latch.

## 6.8.4 Rotational shock, nonoperating

The drive incurs no permanent hardware damage nor data loss when subjected to 30,000 rad/$s^2$ of rotational shock. The shock input is applied around the axis of the actuator pivot for the duration of 1 ms and does not displace the heads from the actuator latch.

# 6.9 Acoustics

The A-weighted sound power levels are given in Bel relative to one picowatt as shown in Figure 44. The measurement method is in accordance with ISO7779. Drives must meet this criteria in the board down orientation.

**36-GB models**

| Mode | Typical (Bels) | Declared (Bels) |
|---|---|---|
| Idle | 4.0 | 4.4 |
| Operating | 4.8 | 5.2 |

**18-GB models**

| Mode | Typical (Bels) | Declared (Bels) |
|---|---|---|
| Idle | 3.8 | 4.2 |
| Operating | 4.7 | 5.0 |

Figure 46. A-weighted sound power (Bels)

Background power levels of the acoustic test chamber for each octave band are recorded. Sound power levels are measured with the drive supported wherein the lower surface of the drive located at a height of 25±3 mm from the chamber floor. No sound-absorbing material external to the drive is used. "Declared" sound power as described by ISO 9296 is defined as

$$LwAd = LwAm + (1.5 * St) + (0.564 * (0.2 - St)) \quad Bels$$

*where*

St = standard deviation
LwAm = sound power mean

The acoustical characteristics of the drive are measured under the conditions specified below:

**Idle mode** - Power on, disks spinning, track following, unit ready to receive and respond to control line commands.

**Operating mode** - Continuous random cylinder selection and Seek operation of the actuator with a dwell time at each cylinder. Seek rate for the drive is calculated as

$$Ns = 0.4/(Tt + T1)$$

*where*

Ns = Average Seek rate in Seeks/second
Tt = Published Seek time from random track to another not including rotational latency
T1 = Time for the drive to rotate by half a revolution

## 6.10 Identification labels

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

- A label containing IBM logo, IBM part number, and the statement "Made by IBM Japan Ltd.", or IBM approved equivalent
- A label containing drive model number, manufacturing date, formatted capacity, country of origin or IBM approved equivalent, and UL/CSA/TUV/CE/CTICK logos
- A bar code label containing the drive serial number
- Jumper setting label
- A user-designed label per agreement
- Interface definition mark - SCSI LVD/SE multi-mode

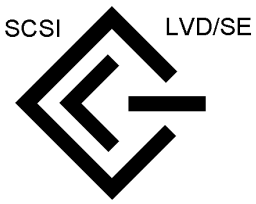


Figure 47. Interface definition mark - SCSI LVD/SE multimode

The labels may be integrated with other labels.

# 6.11 Safety standards (agency approvals)

## 6.11.1 UL and CSA standard conformity

The product is qualified for use in Information Technology Equipment, including Electric Business Equipment, per UL1950 third edition and CAN/CSA C22.2 No. 950-M95 third edition. Either the UL recognition or the CSA certification is maintained for the product life. Either the UL and C-UL recognition mark or the CSA monogram for CSA certification appears on the product.

## 6.11.2 IEC compliance

The product is certified for compliance to IEC 950. The product will comply with these IEC requirements for the life of the product.

## 6.11.3 German safety mark

The product is approved by TUV on Test requirement EN 60950. However, the GS mark is not applicable to internal devices such as these drives.

## 6.11.4 Flammability

Printed circuit boards used in this product are made of material with a UL-recognized flammability rating of V-1 or better. The flammability rating is marked or etched on the board. All other parts not considered electrical components are made of material with a UL-recognized flammability rating of V-1 or better. Small mechanical parts such as cable ties, washers, screws, and PC board mounts are made of material with a UL-recognized flammability rating of V-2.

## 6.11.5 Safe handling

The product is conditioned for safe handling in regards to sharp edges and corners.

## 6.11.6 Environmental precautions

The product does not contain any known or suspected carcinogens. Environmental controls meet or exceed all applicable government regulations in the country of origin. Safe chemical usage and manufacturing control are used to protect the environment. An environmental impact assessment has been done on the manufacturing process used to build the drive, the drive itself, and the disposal of the drive at the end of its life. Production also meets the requirements of the international treaty on chloroflurocarbon (CFC) control known as the United Nations Environment Program Montreal Protocol, and as ratified by the member nations. Materials to be controlled include CFC-11, CFC-12, CFC-113, CFC-114, CFC-115, Halon 1211, Halon 1301, and Halon 2402. Although not specified by the Protocol, CFC-112 is also controlled.

In addition to the Protocol, IBM requires the following:

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

## 6.11.7 Secondary circuit protection

Current limitting devices are provided in termination power and VCM circuit of 12V for overcurrent protection.

# 6.12 Electromagnetic Compatibility

The drive, when installed in a suitable enclosure and exercised with a random accessing routine at maximum data rate, meets the following worldwide EMC requirements:

- United States Federal Communications Commission (FCC) Rules and Regulations (Class B), Part 15. IBM Corporate Standard C-S 2-0001-026
- CISPR Publication 22; 1993 (Class B)
- CISPR Publication 22; Amendment 1, 1995
- CISPR Publication 22; Amendment 2, 1996
- Electrostatic Discharge Susceptibility (ESD) limits for a Class 2 ESD environment specified in IBM Corporate Standard C-S 2-0001-005
- Radiated Electromagnetic Susceptibility (RES) as specified in IBM Corporate Standard C-S 2-0001-012

## 6.12.1 CE mark

The product is declared to be in conformity with requirements of the following EC directives under the sole responsibility of Yamato Lab, IBM Japan Ltd., or IBM United Kingdom Ltd.

Council Directive 89/336/EEC on the approximation of laws of the Member States relating to electromagnetic compatibility.

## 6.12.2 C-TICK mark

The product complies with the Australian EMC standard, AS/NZS 3548:1995 Class B, "Limits and methods of measurement of radio disturbance characteristics of information technology equipment."

This page intentionally left blank.

# Part 2. Interface specification

This page intentionally left blank.

# 7.0 SCSI Command Set

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

| Type | Code | Description |
|---|---|---|
| M | 04h | FORMAT UNIT |
| M | 12h | INQUIRY |
| O | 4Ch | LOG SELECT |
| O | 4Dh | LOG SENSE |
| O | 15h | MODE SELECT (6) |
| O | 55h | MODE SELECT (10) |
| O | 1Ah | MODE SENSE (6) |
| O | 5Ah | MODE SENSE (10) |
| O | 5Eh | PERSISTENT RESERVE IN |
| O | 5Fh | PERSISTENT RESERVE OUT |
| O | 34h | PRE-FETCH |
| M | 08h | READ (6) |
| O | 3Ch | READ BUFFER |
| M | 25h | READ CAPACITY |
| O | 37h | READ DEFECT DATA (10) |
| O | B7h | READ DEFECT DATA (12) |
| M | 28h | READ (10) |
| O | 3Eh | READ LONG |
| O | 07h | REASSIGN BLOCKS |
| O | 1Ch | RECEIVE DIAGNOSTICS |
| M | 17h | RELEASE (6) |
| O | 57h | RELEASE (10) |
| O | A3h | REPORT DEVICE IDENTIFIER |
| O | A0h | REPORT LUNS |
| M | 03h | REQUEST SENSE |
| M | 16h | RESERVE (6) |
| O | 56h | RESERVE (10) |
| O | 01h | REZERO UNIT |
| O | 0Bh | SEEK (6) |
| O | 2Bh | SEEK (10) |
| M | 1Dh | SEND DIAGNOSTIC |
| O | 1Bh | START STOP UNIT |
| O | 35h | SYNCHRONIZE CACHE |
| M | 00h | TEST UNIT READY |
| O | 2Fh | VERIFY |
| M | 0Ah | WRITE (6) |
| O | 2Eh | WRITE AND VERIFY |
| O | 3Bh | WRITE BUFFER |
| M | 2Ah | WRITE (10) |
| O | 3Fh | WRITE LONG |
| O | 41h | WRITE SAME |

Figure 48. Supported SCSI commands

## 7.1 SCSI Control Byte

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

| Bit | | | | | | | |
|---|---|---|---|---|---|---|---|
| **7** | **6** | **5** | **4** | **3** | **2** | **1** | **0** |
| VU = 0 | | Reserved = 0 | | | | FLAG | LINK |

Figure 49. SCSI Control Byte

**VU** VU stands for Vendor Unique.

**FLAG** The Flag bit specifies which message the target shall return to the initiator if the link bit is one and the command completes without any error. If Link is zero, Flag must also be zero. If Link is one and the command terminates successfully, the drive will send either the LINKED TASK COMPLETE message (FLAG=0) or the LINKED TASK COMPLETE WITH FLAG message (FLAG=1). Typically this bit is used to cause an interrupt in the initiator between linked commands.

**LINK** This bit is set to one to indicate that the initiator desires an automatic link to the next command upon successful completion of the current command. Upon successful completion of the command, the drive will return INTERMEDIATE GOOD status and then send one of the two messages defined under Flag above.

Upon unsuccessful completion of the command, the drive will return CHECK CONDITION status or RESERVATION CONFLICT status and then send the TASK COMPLETE message. No further commands in the chain are executed.

## 7.2 Abbreviations

These abbreviations are used throughout the following sections.

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

**VU** Vendor Unique bits

**LBA** Logical Block Address

**RSVD** Reserved

**MSB** Most Significant bit

**LSB** Least Significant bit

## 7.3 Byte ordering conventions

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

# 7.4 FORMAT UNIT (04)

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
| 0 | Command Code = 04h | | | | | | | |
| 1 | Reserved | | | FmtData | CmpList | Defect List Format | | |
| 2 | VU = 0 | | | | | | | |
| 3<br>4 | (MSB) Interleave Factor (LSB) | | | | | | | |
| 5 | VU = 0 | | Reserved = 0 | | | | FLAG | LINK |

Figure 50. FORMAT UNIT (04)

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

- **FmtData** set to one specifies that a Data Out phase follows the Command phase. FmtData set to zero specifies that no Data Out phase follows.
- **CmpList** set to one specifies that the GList (Grown Defect List) existing prior to the format **not** be used and is discarded. The drive is formatted with PList and DList (if specified). DList becomes the new GList.

  **Note:** The drive manages two internal defect lists and one external. The primary defect list (“P”List) is created at time of manufacture. The grown defect list (“G”List) is built after manufacture by the Initiators use of the REASSIGN BLOCK command and the Automatic Reallocate functions. The data defect list (“D”List) is an external list. It is supplied by the initiator in the DATA OUT phase of the FORMAT UNIT command.
- **Defect List Format** specifies the format of the defect descriptor transferred to the Target when FmtData bit is set to one. The Target supports three defect descriptor formats for the Format Unit command as following:

  **Format Description**

  **000b** Block format

  **100b** Bytes From Index format

  **101b** Physical Sector format

  If the FmtData bit is set to zero, this field must also be zero. Otherwise the command will complete with a check condition with a sense key of illegal request and an additional sense code of invalid field in CDB.

  **Interleave Factor** may be zero or one, either of which specifies an interleave of 1:1. Other Interleave Factors are ignored because of the extensive buffering implemented in the drive.

## 7.4.1 Defect list

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

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
| 0 | Reserved = 0 | | | | | | | |
| 1 | FOV | DPRY | DCRT | STPF = 1 | IP = 0 | DSP= 0 | Immed | 0 |
| 2<br>3 | (MSB) Defect List Length (LSB) | | | | | | | |

Figure 51. Format of Defect List Header

The Target has a limited implementation of the Format Option bits located in Bits 2 through 7 of Byte 1 of the Defect List Header (See Figure 48). If the Initiator attempts to select any function not implemented by the Target, the Target terminates the command with Check Condition Status. The sense key is set to Illegal Request and the additional sense code is set to Invalid Field in Parameter List.

- **FOV** (Format Options Valid) bit of zero causes the Target to verify that the setting for the DPRY (Disable Primary), DCRT (Disable Certification), STPF (Stop Format), IP (Initialize Pattern), and DSP (Disable Saving Parameters) bits are zero. If any of these bits are not zero, the Target terminates the command with Check Condition Status. The sense key is set to Illegal Request and the additional sense code is set to Invalid Field in Parameter List.

  **Note:** When FOV bit is one, three combinations of the DPRY, DCRT, STPF, IP and DSP bits are allowed. Any other combinations return a Check Condition Status With a sense key of Illegal Request and an additional sense code of Invalid Field In Parameter List. The supported combination are

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

- **DPRY** (Disable Primary) bit set to zero indicates that the Target does not use portions of the medium identified as defective in the primary defect PList for Initiator addressable logical blocks. If the Target cannot locate the PList or it cannot determine whether a PList exists, the target terminates the Format Unit command as described for STPF=1. A DPRY bit of one indicates that the Target does not use the Plist to identify defective areas of the medium. The Plist is not deleted.
- **DCRT** (Disable Certification) bit of ZERO indicates that the Target performs a medium certification operation and generates a Certification List (CList) and the Target adds the Clist to the Glist. A DCRT bit of one indicates that the Target does not generate a CList (Certification List) nor perform a certification process while executing the Format Unit Command.

  **Note:** Since the DCRT bit is part of the Data Out phase that follows the format command, the FCERT bit in Mode Page 0 is provided to control certification when the format command is issued with no Data Out phase. If a format command is issued with a Data Out phase then FCERT is ignored.

- **STPF** (Stop Format) bit must be set to one. If one or both of the following conditions occurs, the Target terminates the Format Unit command with Check Condition Status. The sense key is set to Medium Error and the additional sense code is set to either Defect List Not Found if the first condition occurred or it is set to Defect List Error if the second condition occurred.

  1. The Target cannot locate a required defect list nor determine that the list exists.
  2. The Target encounters an unrecoverable error while accessing a required defect list.

- **IP** (Initialization Pattern) bit must be set to zero. The Target initializes all data with zeros.
- **DSP** (Disable Saving Parameters) bit must be set to zero. The Target saves all the Mode Select savable parameters during the format operation.
- **Immed** (Immediate) bit set to zero requests that status be returned at the end of the format operation. An immediate bit set to one requests that status be returned immediately. Good Status is returned following the CDB validation and transfer of data in the Data Out phase. If the immediate format operation terminates in error, Deferred Error Sense data is generated. With the immediate bit set to one, the Link bit must be set to zero.

## 7.4.2 Defect descriptor

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

The Target supports three Defect List formats.

### 7.4.2.1 Block format

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

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
| 0<br>1<br>2<br>3 | (MSB)<br>Defective Logical Block Address<br>(LSB) | | | | | | | |
| 4n -<br>4n +3 | Defective Logical Block Address n | | | | | | | |

Figure 52. Defect descriptor - Block format.

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

### 7.4.2.2 Bytes From Index format

Each defect descriptor for the Bytes From Index format specifies that the sector containing this byte be marked defective. The defect descriptor is comprised of the cylinder number of the defect, the head number of the defect, and the defect bytes from index.

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
| 0<br>1<br>2 | (MSB)<br>Cylinder Number of Defect<br>(LSB) | | | | | | | |
| 3 | Head Number of Defect | | | | | | | |
| 4<br>5<br>6<br>7 | (MSB)<br>Defect Bytes from Index<br>(LSB) | | | | | | | |
| 8n -<br>8n +7 | Defect Descriptor n | | | | | | | |

Figure 53. Defect descriptor - Bytes from Index format.

Format of the defect list sent during the data out phase when FmtData is set to one.

### 7.4.2.3 Physical Sector format

Each defect descriptor for the Physical Sector format specifies a defect that is the length of a sector. The defect descriptor is comprised of the cylinder number of the defect, the head number of the defect, and the defect sector number.

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
| 0<br>1<br>2 | (MSB)<br>Cylinder Number of Defect<br>(LSB) | | | | | | | |
| 3 | Head Number of Defect | | | | | | | |
| 4<br>5<br>6<br>7 | (MSB)<br>Defect Sector Number<br>(LSB) | | | | | | | |
| 8n -<br>8n + 7 | Defect Descriptor n | | | | | | | |

Figure 54. Defect descriptor - Physical Sector format.

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

## 7.4.3 INQUIRY (12)

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | **7** | **6** | **5** | **4** | **3** | **2** | **1** | **0** |
| 0 | Command Code = 12h | | | | | | | |
| 1 | Reserved | | | Reserved = 0 | | | CmdDt | EVPD |
| 2 | Page Code | | | | | | | |
| 3 | Reserved = 0 | | | | | | | |
| 4 | Allocation Length | | | | | | | |
| 5 | VU =0 | | Reserved = 0 | | | | FLAG | LINK |

Figure 55. INQUIRY (12)

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

**A CmdDT bit of one** specifies that the drive shall return the command support data information identified by the Page Code field in the CDB.

**An EVPD bit of one** specifies that the drive return the vital product data page identified by the Page Code field in the CDB[1]. **Page code** specifies which page of vital product data information the drive shall return.

[1] The available VPD pages are defined in the addendum provided for each different drive model in the section entitled *Inquiry Data Format.*

| CmdDt | EVPD | PAGE CODE | Description |
|---|---|---|---|
| 0 | 0 | 0 | The drive returns the standard INQUIRY data. |
| 0 | 0 | Non Zero | The drive returns CHECK CONDITION status with the sense key of ILLEGAL REQUEST and the additional sense code of INVALID FIELD IN CDB. |
| 0 | 1 | Supported | The drive returns the vital product data of page code requested. |
| 0 | 1 | Unsupported | The drive returns CHECK CONDITION status with the sense key of ILLEGAL REQUEST and the additional sense code of INVALID FIELD IN CDB |
| 1 | 0 | Supported | The drive returns command support data of page code requested. |
| 1 | 0 | Unsupported | The drive returns CHECK CONDITION status with the sense key of ILLEGAL REQUEST and the additional sense code of INVALID FIELD IN CDB. |
| 1 | 1 | Any values | The drive returns CHECK CONDITION status with the sense key of ILLEGAL REQUEST and the additional sense code of INVALID FIELD IN CDB. |

Figure 56. Page code descriptions

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

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

## 7.4.4 INQUIRY DATA

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

### 7.4.4.1 Inquiry data format - CmdDt = 0 EVPD = 0

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | **7** | **6** | **5** | **4** | **3** | **2** | **1** | **0** |
| 0 | Qualifier = 0 | | | Peripheral Device Type = 0 | | | | |
| 1 | RMB=0 | Device-Type Modifier = 0 | | | | | | |
| 2 | ISO = 0 | | ECMA = 0 | | | ANSI = 3 | | |
| 3 | RSVD = 0 | TrmTsk = 0 | Norm ACA=0 | RSVD = 0 | Response Data Format = 2 | | | |
| 4 | Additional Length = 159 (9Fh) | | | | | | | |
| 5 | Reserved = 0 | | | | | | | |
| 6 | RSVD = 0 | EncSer = 0 | RSVD = 0 | MultiP =0 | MChngr = 0 | ACKREQ = 0 | Addr32 = 0 | Addr16 = 1 |
| 7 | REL_A = 0 | Wb_32 = 0 | Wb_16 = 1 | Sync = 1 | Link = 1 | TTD = 0 | CmdQu = 1 | RSVD = 0 |
| 8-15 | Vendor ID = 'IBM ' (ASCII) | | | | | | | |
| 16-31 | Product ID (ASCII) | | | | | | | |
| 32-35 | Product Revision Level (ASCII) | | | | | | | |
| 36-43 | Unit Serial Number (ASCII) | | | | | | | |
| 44-55 | Reserved = 0 | | | | | | | |
| 56 | Reserved = 0 | | | | Clocking=11b | | QAS=0 | IUS=0 |
| 57-95 | Reserved = 0 | | | | | | | |
| 96-145 | Copyright Notice (ASCII) | | | | | | | |
| 146-163 | Reserved = 0 | | | | | | | |

Figure 57. INQUIRY Data - CmdDt = 0EVPD = 0

- **Qualifier** is set to zero to indicate that the LUN specified in the Identify Message is currently supported. Qualifier is set to 011b when the LUN specified in the Identify Message is not present.[2]
- **Peripheral Device Type** is set to zero to indicate that the device is Direct Access. Peripheral Dev. Type is set to 1Fh when the LUN specified in the Identify Message is not present.
- **Removal Media Bit (RMB)** is always set to zero to indicate no removal media exist.
- **Device-Type Modifier** is set to zero.

[2] For all commands, except inquiry and request sense, if an invalid lun is specified a check condition will be returned.

- **ISO** is set to zero to indicate that this product does not claim compliance to the International Organization for Standardization (ISO) version of SCSI (ISO DP 9316).
- **ECMA** is set to zero to indicate that this product does not claim compliance to the European Computer Manufacturers Association (ECMA) version of SCSI (ECMA-111).
- **ANSI** indicates the level of the ANSI standard that is supported by the product. The drive supports ANSI SCSI version 3.
- **TrmTsk** (Terminate Task) filed of 0 indicates the drive does not support the TERMINATE TASK task management function as defined in the SAM.
- **NormACA** (Normal ACA) field of 0 indicates the device server does not support setting the NACA bit to one in the Control Byte of the CDB as defined in the SAM.
- **Response Data Format** is set to two to indicate that the Inquiry Data Format as specified in ANSI SCSI version 2 is supported by the drive.
- **Additional Length** indicates the number of bytes of inquiry information that follows.
- **EncSer** (Enclosure Services) bit of 0 indicates that the drive does not contain an embedded enclosure service component.
- **MultiP** (MultiPort) bit of 0 indicates that the drive has a single port and does not implement multi- port requirements.
- **MChngr** (Medium Changer) bit is always 0 to indicate MChngr is not supported.
- **ACKREQQ** (ACKQ/REQQ) bit of 0 indicates that the drive does not support a request and acknowledge data transfer handshake on a Q cable.
- **Addr32** (Wide SCSI Address 32) bit of 0 indicates that the drive does not support 32 bit wide SCSI Addresses.
- **Addr16** (Wide SCSI Address 16) bit of 1 indicates that the drive supports 16 bit wide SCSI Addresses.
- **REL_A** is set to zero to indicate that the drive does not support 'Relative Address Mode'.
- **Wb_32** is set to zero to indicate that the drive does not support 32-bit wide data transfers.
- **Wb_16** is set to one to indicate that the drive supports 16-bit wide data transfers.
- **Sync** is set to one to indicate that the drive supports synchronous data transfer.
- **Link** is set to one to indicate that the drive supports linked commands.
- **TTD** is set to zero to indicate that the drive does not support the CONTINUE I/O PROCESS and TARGET TRANSFER DISABLE message for this logical unit.
- **CmdQu** is set to one to indicate that the drive supports command queuing.
- **Vendor ID** is 'IBM' padded with ASCII blanks.
- **Product ID** is specified in ASCII character as shown:

| Product ID | Description |
|---|---|
| IC35L018UW | 15000RPM 18.3GB 68pin |
| IC35L018UC | 15000RPM 18.3GB 80pin |
| IC35L036UW | 15000RPM 36.7GB 68pin |
| IC35L036UC | 15000RPM 36.7GB 80pin |

- **Product Revision Level** indicates the level of microcode. It indicates RAM microcode level.
- **Unit Serial Number** contains the drive serial number.

- **Clocking** 11b indicates the target supports ST (Single Transition) and DT (Double Transition) on synchronous transfer.
- **QAS** (Quick Arbitrate Supported) bit of zero indicates that the target does not support the quick arbitrate feature.
- **IUS** (Information Unit Supported) bit of zero indicates that the target does not support information unit.

### 7.4.4.2 Inquire data format - CmdDt = 1

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | **7** | **6** | **5** | **4** | **3** | **2** | **1** | **0** |
| 0 | Peripheral Qualifier | | | Peripheral Device Type | | | | |
| 1 | Reserved | | | | | Support | | |
| 2 | Version | | | | | | | |
| 3 | Reserved | | | | | | | |
| 4 | Reserved | | | | | | | |
| 5 | CDB Size = m-5 | | | | | | | |
| 6 | CDB Usage Data | | | | | | | |
| m | | | | | | | | |

Figure 58. Command support data format

**Peripheral Qualifier** is set to zero to indicate that the LUN specified in the Identify Message is currently supported. Peripheral Qualifier is set to 011b when the LUN specified in the Identify Message is not present.[3]

- **Peripheral Device Type** is set to zero to indicate that the device is a Direct-Access. Peripheral Device Type is set to 1Fh when the LUN specified in the Identify Message is not present.
- **Support** indicates the type of command support. The following table defines the values and meanings of the Support field.

[3] For all commands, except inquiry and request sense, if an invalid LUN is specified a check condition will be returned.

| Support | Description |
|---|---|
| 000b | Data about the requested SCSI operation code is not currently available. |
| 001b | The device server does not support the tested SCSI operation code. All data after byte 1 is undefined. |
| 010b | Reserved. |
| 011b | The device server supports the tested SCSI operation code in conformance with the SCSI standard. |
| 100b | Vendor-specific |
| 101b | The device server supports the tested SCSI operation code in a vendor-specific manner. |
| 110b | Vendor-specific |
| 111b | Reserved |

Figure 59. SUPPORT values and meanings

**Version** indicates the level of the ANSI standard that is supported by the product. The drive supports ANSI SCSI version 3.

- **CDB Size** indicates the number of bytes in the CDB for the operation, code, and size of the CDB Usage Data field in the return data.
- **CDB Usage Data** contains information about the CDB for the operation code. The first byte of the CDB Usage Data contains the operation code. All bytes except the first byte of the CDB Usage Data contain a usage map for bits in the CDB. In the usage map, the bit that is all or part of a field in the CDB is set to one, otherwise it is set to zero.

### 7.4.4.3 Inquiry data format - EVPD = 1 - Page Code = 00

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
| 0 | Qualifier = 0 | | | Peripheral Device Type = 0 | | | | |
| 1 | Page Code = 00h | | | | | | | |
| 2 | Reserved = 0 | | | | | | | |
| 3 | Page Length = 03h | | | | | | | |
| 4 | Supported Page Code = 00h | | | | | | | |
| 5 | Supported Page Code = 80h | | | | | | | |
| 6 | Supported Page Code = 83h | | | | | | | |

Figure 60. INQUIRY DATA - EVPD = 1 (Page Code = 00)

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

### 7.4.4.4 Inquiry data format - EVPD = 1 - Page Code = 80h

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

Figure 61. INQUIRY DATA - EVPD = 1 (Page Code = 80h)

- **Qualifier** is set to zero to indicate that the LUN specified in the Command Block is currently supported.
- **Peripheral Device Type** is set to zero to indicate that the device is Direct Access.
- **Page Code** is set to 80h, and this field contains the same value as in the page code field of the INQUIRY command descriptor block.
- **Page length** is set to 16, and this field specifies the length of the following page data.
- **Serial Number** gives the drive serial number right aligned.

### 7.4.4.5 Inquiry data format - EVPD = 1 - Page Code = 83h

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

Figure 62. INQUIRY DATA - EVPD = 1 (Page Code = 83h)

- **Qualifier** is set to zero to indicate that the LUN specified in the Command Block is currently supported.
- **Peripheral Device Type** is set to zero to indicate that the device is Direct Access.
- **Page Code** is set to 83h, and this field contains the same value as in the page code field of the INQUIRY command descriptor block.
- **Page length** is set to 12, and this field specifies the length of the following page data.
- **Code Set** field specifies the code set used for the identifier field. The Target supports binary.
- **Association** field specifies the entity with which the identifier field is associated. The Target supports value of 0h, the Identifer field is associated with the addressed physical or logical device.
- **Identifier Type** field specifies the format and assignment authority for the identifier. The Target supports the value of 03h.
- **World Wide ID** is 64 bit unique identification for each drive. The format is: **5005076**h **xxx**h **yy**b **n**. where
  - **xxx** is 12 bit Block assignment defined for each model and manufacturing site.
  - **yy** is 2 bit Port/Node ID select. This is 11b as for parallel SCSI interface devices.
  - **n** is 22 bit drive unique serial number.

# 7.5 LOG SELECT (4C)

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

Figure 63. LOG SELECT (4C)

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

- **PCR** The Parameter Code Reset determines whether the Log Sense parameters will be cleared and unit attention posted for all other initiators. A value of 1 indicates that the parameters be cleared, while a value of zero (except when PC = 11b) indicates that the parameters not be cleared.
- **SP** The Save Parameters bit value of zero indicates that the page parameters not be saved. A value of 1 indicates that the page parameters that are savable be saved after they have been changed.
- **PC** The Page Control field defines the type of parameters to be selected. The PC field set to 11b (and PCR is then a don't care) indicates that the Default Cumulative values are set to their default values of 0. If the PC field is set to 01b and PCR is set to 1, the Current Cumulative values are also set to their default values.
- **Parameter List Length** The Parameter List Length field specifies the length in bytes of the parameter list that shall be located in the DATA OUT buffer. A parameter list length zero indicates that no pages shall be transferred.

If one or more fields of the CDB are not set correctly the command will be terminated with a CHECK CONDITION status. The Sense Key shall be set to Illegal Request and the additional sense code set to Invalid Field in CDB.

The Log Select command will reset the counter variables to their default values of zero. These variables are listed in the Log Sense command.

To indicate that parameters have changed, the target generates a unit attention condition for all initiators except the one that issued the Log Select command.

# 7.6 LOG SENSE (4D)

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

Figure 64. LOG SENSE (4D)

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

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

## 7.6.1 Log Page parameters

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

- **Page Header**
  Page Code field identifies which log page is being transferred.

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

- **Log Parameters**
  Each log parameter begins with a four-byte parameter header followed by one or more bytes of parameter value data.

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

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

  - **DU** The Disable Update bit is set to 0 which indicates that the drive updates the log parameter value to reflect events that should be noted by that parameter.
  - **DS** The Disable Save bit is set to 1 to indicate the parameter is non-savable and is set to 0 to indicate that the parameter is savable.
  - **TSD** The Target Save Disable bit is set to zero which indicates that the drive provides a target defined method for saving log parameters.
  - **ETC** The enable Threshold Comparison bit is set to 0 which indicates the drive does not perform comparisons between cumulative and any threshold values.
  - **TMC** The Threshold Met Criteria field is not valid because this drive does not perform threshold comparisons. This field is set to 0.
  - **LBIN** The List Binary bit is only valid if the LP is 1. If the LP bit is 1 and the LBIN bit is 0, then the list parameter is a string of ASCII graphic code (20h - 7Eh). If the LP bit is 1 and the LBIN bit is 1, then the list parameter is a list of binary information.
  - **LP** The List Parameter bit is set to 0 for parameters that are data counters. The LP bit is set to 1 for parameters that are lists.

## 7.6.2 Log Sense Page 0

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

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
| 0 | Reserved | | Page code = 0 | | | | | |
| 1 | Reserved | | | | | | | |
| 2-3 | Page Length = 000Ah (Number of Pages Supported) | | | | | | | |
| 4 | First supported page 0h | | | | | | | |
| 5 | Second supported page 2h | | | | | | | |
| 6 | Third supported page 3h | | | | | | | |
| 7 | Fourth supported page 5h | | | | | | | |
| 8 | Fifth supported page 6h | | | | | | | |
| 9 | Sixth supported page Dh | | | | | | | |
| 10 | Seventh supported page Eh | | | | | | | |
| 11 | Eighth supported page Fh | | | | | | | |
| 12 | Ninth supported page 10h | | | | | | | |
| 13 | Tenth supported page 2Fh | | | | | | | |

Figure 65. Log Sense Page 0

## 7.6.3 Log Sense Page 2

This page contains counters for write errors.

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | **7** | **6** | **5** | **4** | **3** | **2** | **1** | **0** |
| 0 | Reserved | | Page code = 02h | | | | | |
| 1 | Reserved | | | | | | | |
| 2-3 | PageLength = 54h | | | | | | | |
| 4-5 | Parameter Code = 00h | | | | | | | |
| 6 | DU = 0 | DS = 0 | TSD = 0 | ETC = 0 | TMC = 0 | | LBIN = 0 | LP = 0 |
| 7 | Parameter Length = 08h | | | | | | | |
| 8-15 | Errors recovered without delay = 0 | | | | | | | |
| 16-17 | Parameter Code = 01h | | | | | | | |
| 18 | DU = 0 | DS = 0 | TSD = 0 | ETC = 0 | TMC = 0 | | LBIN = 0 | LP = 0 |
| 19 | Parameter Length = 08h | | | | | | | |
| 20-27 | Count of LBA's with write fault errors | | | | | | | |
| 28-29 | Parameter Code = 02h | | | | | | | |
| 30 | DU = 0 | DS = 0 | TSD = 0 | ETC = 0 | TMC = 0 | | LBIN = 0 | LP = 0 |
| 31 | Parameter Length = 08h | | | | | | | |
| 32-39 | Count of LBA's with id type errors | | | | | | | |
| 40-41 | Parameter Code = 03h | | | | | | | |
| 42 | DU = 0 | DS = 0 | TSD = 0 | ETC = 0 | TMC = 0 | | LBIN = 0 | LP = 0 |
| 43 | Parameter Length = 08h | | | | | | | |
| 44-51 | Total errors recovered | | | | | | | |
| 52-53 | Parameter Code = 04h | | | | | | | |
| 54 | DU = 0 | DS = 0 | TSD = 0 | ETC = 0 | TMC = 0 | | LBIN = 0 | LP = 0 |
| 55 | Parameter Length = 08h | | | | | | | |
| 56-63 | Times recovery invoked | | | | | | | |

Figure 66. Log Sense Page 2 (part 1 of 2)

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | **7** | **6** | **5** | **4** | **3** | **2** | **1** | **0** |
| 64-65 | Parameter Code = 05h | | | | | | | |
| 66 | DU = 0 | DS = 0 | TSD = 0 | ETC = 0 | TMC = 0 | | LBIN = 0 | LP = 0 |
| 67 | Parameter Length = 08h | | | | | | | |
| 68-75 | Total bytes written | | | | | | | |
| 76-77 | Parameter Code = 06h | | | | | | | |
| 78 | DU = 0 | DS = 0 | TSD = 0 | ETC = 0 | TMC = 0 | | LBIN = 0 | LP = 0 |
| 79 | Parameter Length = 08h | | | | | | | |
| 80-87 | Count of LBA's with hard errors | | | | | | | |

Figure 67. Log Sense Page 2 (part 2 of 2)

## 7.6.4 Log Sense Page 3

This page contains counters for read errors

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | **7** | **6** | **5** | **4** | **3** | **2** | **1** | **0** |
| 0 | Reserved | | Page code = 03h | | | | | |
| 1 | Reserved | | | | | | | |
| 2-3 | PageLength = 54h | | | | | | | |
| 4-5 | Parameter Code = 00h | | | | | | | |
| 6 | DU = 0 | DS = 0 | TSD = 0 | ETC = 0 | TMC = 0 | | LBIN = 0 | LP = 0 |
| 7 | Parameter Length = 08h | | | | | | | |
| 8-15 | Errors recovered without delay = 0 | | | | | | | |
| 16-17 | Parameter Code = 01h | | | | | | | |
| 18 | DU = 0 | DS = 0 | TSD = 0 | ETC = 0 | TMC = 0 | | LBIN = 0 | LP = 0 |
| 19 | Parameter Length = 08h | | | | | | | |
| 20-27 | Count of LBA's with ECC detected errors | | | | | | | |
| 28-29 | Parameter Code = 02h | | | | | | | |
| 30 | DU = 0 | DS = 0 | TSD = 0 | ETC = 0 | TMC = 0 | | LBIN = 0 | LP = 0 |
| 31 | Parameter Length = 08h | | | | | | | |
| 32-39 | Count of LBA's with id type errors | | | | | | | |
| 40-41 | Parameter Code = 03h | | | | | | | |
| 42 | DU = 0 | DS = 0 | TSD = 0 | ETC = 0 | TMC = 0 | | LBIN = 0 | LP = 0 |
| 43 | Parameter Length = 08h | | | | | | | |
| 44-51 | Total errors recovered | | | | | | | |
| 52-53 | Parameter Code = 04h | | | | | | | |
| 54 | DU = 0 | DS = 0 | TSD = 0 | ETC = 0 | TMC = 0 | | LBIN = 0 | LP = 0 |
| 55 | Parameter Length = 08h | | | | | | | |
| 56-63 | Times recovery invoked | | | | | | | |
| 64-65 | Parameter Code = 05h | | | | | | | |

Figure 68. Log Sense Page 3 (part 1 of 2)

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | **7** | **6** | **5** | **4** | **3** | **2** | **1** | **0** |
| 66 | DU = 0 | DS = 0 | TSD = 0 | ETC = 0 | TMC = 0 | | LBIN = 0 | LP = 0 |
| 67 | Parameter Length = 08h | | | | | | | |
| 68-75 | Total bytes written | | | | | | | |
| 76-77 | Parameter Code = 06h | | | | | | | |
| 78 | DU = 0 | DS = 0 | TSD = 0 | ETC = 0 | TMC = 0 | | LBIN = 0 | LP = 0 |
| 79 | Parameter Length = 08h | | | | | | | |
| 80-87 | Count of LBA's with hard errors | | | | | | | |

Figure 69. Log Sense Page 3 (part 2 of 2)

The drive will attempt to read data after a seek before the head has fully settled on track. This is done to aid performance. However, as a result there is an increased incidence of error recovery invoked which normally uses ECC or retries to recover the data. As a consequence of this an error recovered by a single retry is not reported by the error counters.

Additionally the drive does not report data recovered by ECC on the fly.

## 7.6.5 Log Sense Page 5

This page contains counters for verify errors

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | **7** | **6** | **5** | **4** | **3** | **2** | **1** | **0** |
| 0 | Reserved | | Page code = 05h | | | | | |
| 1 | Reserved | | | | | | | |
| 2-3 | PageLength = 54h | | | | | | | |
| 4-5 | Parameter Code = 00h | | | | | | | |
| 6 | DU = 0 | DS = 0 | TSD = 0 | ETC = 0 | TMC = 0 | | LBIN = 0 | LP = 0 |
| 7 | Parameter Length = 08h | | | | | | | |
| 8-15 | Errors recovered without delay = 0 | | | | | | | |
| 16-17 | Parameter Code = 01h | | | | | | | |
| 18 | DU = 0 | DS = 0 | TSD = 0 | ETC = 0 | TMC = 0 | | LBIN = 0 | LP = 0 |
| 19 | Parameter Length = 08h | | | | | | | |
| 20-27 | Count of LBA's with ECC detected errors | | | | | | | |
| 28-29 | Parameter Code = 02h | | | | | | | |
| 30 | DU = 0 | DS = 0 | TSD = 0 | ETC = 0 | TMC = 0 | | LBIN = 0 | LP = 0 |
| 31 | Parameter Length = 08h | | | | | | | |
| 32-39 | Count of LBA's with id type errors | | | | | | | |
| 40-41 | Parameter Code = 03h | | | | | | | |
| 42 | DU = 0 | DS = 0 | TSD = 0 | ETC = 0 | TMC = 0 | | LBIN = 0 | LP = 0 |
| 43 | Parameter Length = 08h | | | | | | | |
| 44-51 | Total errors recovered | | | | | | | |
| 52-53 | Parameter Code = 04h | | | | | | | |
| 54 | DU = 0 | DS = 0 | TSD = 0 | ETC = 0 | TMC = 0 | | LBIN = 0 | LP = 0 |
| 55 | Parameter Length = 08h | | | | | | | |
| 56-63 | Times recovery invoked | | | | | | | |
| 64-65 | Parameter Code = 05h | | | | | | | |

Figure 70. Log Sense Page 5 (part 1 of 2)

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | **7** | **6** | **5** | **4** | **3** | **2** | **1** | **0** |
| 66 | DU = 0 | DS = 0 | TSD = 0 | ETC = 0 | TMC = 0 | | LBIN = 0 | LP= 0 |
| 67 | Parameter Length = 08h | | | | | | | |
| 68-75 | Total Bytes Verified | | | | | | | |
| 76-77 | Parameter Code = 06h | | | | | | | |
| 78 | DU = 0 | DS = 0 | | TSD = 0 | TMC = 0 | | LBIN = 0 | LP = 0 |
| 79 | Parameter Length = 08h | | | | | | | |
| 80-87 | Count of LBA's with hard errors | | | | | | | |

Figure 71. Log Sense Page 5 (part 2 of 2)

In order to aid performance, the drive will attempt to read data after a seek before the head has fully settled on track. However there is an increased incidence of error recovery invoked as a result which normally uses ECC or a retry to recover the data. As a consequence of this an error recovered by a single retry is not reported by the error counters.

Additionally the drive does not report data recovered by ECC on the fly.

## 7.6.6 Log Sense Page 6

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

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

Figure 72. Log Sense Page 6

## 7.6.7 Log Sense Page D

This page contains temperature information.

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

Figure 73. Log Sense Page D

## 7.6.8 Log Sense Page E

This page contains the start-stop cycle information.

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | **7** | **6** | **5** | **4** | **3** | **2** | **1** | **0** |
| 0 | Reserved | | Page code = 0Eh | | | | | |
| 1 | Reserved | | | | | | | |
| 2-3 | PageLength = 24h | | | | | | | |
| 4-5 | Parameter Code = 0001h | | | | | | | |
| 6 | DU = 0 | DS = 1 | TSD = 0 | ETC = 0 | TMC = 0 | | LBIN = 0 | LP = 1 |
| 7 | Parameter Length = 06h | | | | | | | |
| 8-11 | Year of Manufacture (4 ASCII characters) | | | | | | | |
| 12-13 | Week of Manufacture (2 ASCII characters) | | | | | | | |
| 14-15 | Parameter Code 0002h | | | | | | | |
| 16 | DU = 0 | DS = 1 | TSD = 0 | ETC = 0 | TMC = 0 | | LBIN = 0 | LP = 1 |
| 17 | Parameter Length = 06h | | | | | | | |
| 18-21 | Accounting Date Year (4 ASCII characters) | | | | | | | |
| 22-23 | Accounting Date Week (2 ASCII characters) | | | | | | | |
| 24-25 | Parameter Code 0003h | | | | | | | |
| 26 | DU = 0 | DS = 1 | TSD = 0 | ETC = 0 | TMC = 0 | | LBIN = 1 | LP = 1 |
| 27 | Parameter Length = 04h | | | | | | | |
| 28-31 | Specified cycle count over device lifetime | | | | | | | |
| 32-33 | Parameter Code 0004h | | | | | | | |
| 34 | DU = 0 | DS = 1 | TSD = 0 | ETC = 0 | TMC = 0 | | LBIN = 1 | LP = 1 |
| 35 | Parameter Length = 04h | | | | | | | |
| 36-39 | Accumulated start-stop cycles (4 byte binary number) | | | | | | | |

Figure 74. Log Sense Page E

The week and year that the device was manufactured shall be set in the parameter field defined by parameter code 0001h. The date of manufacture cannot be saved using the LOG SELECT command. The data is expected in numeric ASCII characters (30-39h) in the form YYYYWW. The accounting date specified by parameter code 0002h is a parameter that can be saved using the LOG SELECT command.

## 7.6.9 Log Sense Page F

This page contains the Application Client Log.

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | **7** | **6** | **5** | **4** | **3** | **2** | **1** | **0** |
| 0 | Reserved | | Page code = 0Fh | | | | | |
| 1 | Reserved | | | | | | | |
| 2-3 | Page length = 4000h | | | | | | | |
| | Application client log parameter | | | | | | | |
| 4- | First application client log parameter | | | | | | | |
| | | | | | | | | |
| -4003h | Last application client log parameter | | | | | | | |

The following table describes the application client log parameter structure.

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | **7** | **6** | **5** | **4** | **3** | **2** | **1** | **0** |
| 0-1 | Parameter code | | | | | | | |
| 2 | DU = 1 | DS = 0 | TSD = 0 | ETC = 0 | TMC = 0 | | LBIN = 1 | LP = 1 |
| 3 | Parameter length = FCh | | | | | | | |
| 4- | First parameter byte | | | | | | | |
| | | | | | | | | |
| 255 | Last parameter byte | | | | | | | |

Figure 75. Log Sense Page F Application Client Log

Parameter code 0000h through 003Fh are supported.

The values stored in the parameter bytes represent data sent to the device in a previous LOG SELECT command.

## 7.6.10 Log Sense Page 10

This page contains Self-test results. The results of the twenty most recent self-tests are stored in this Log page.

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
| 0 | Reserved | | Page code = 10h | | | | | |
| 1 | Reserved | | | | | | | |
| 2-3 | PageLength = 190h | | | | | | | |
| | Self-test results log parameters | | | | | | | |
| 4- | First self-test results log parameter | | | | | | | |
| | | | | | | | | |
| -403h | Last self-test results log parameter | | | | | | | |

The following table describes the self-test results log parameter structure.

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
| 0-1 | Parameter code | | | | | | | |
| 2 | DU = 0 | DS = 0 | TSD = 0 | ETC = 0 | TMC = 0 | | LBIN = 1 | LP = 1 |
| 3 | Parameter Length = 10h | | | | | | | |
| 4 | Function Code | | | RSVD | Self-Test Results Value | | | |
| 5 | Extended Segment Number | | | | | | | |
| 6-7 | Timestamp | | | | | | | |
| 8-15 | LBA of First Failure | | | | | | | |
| 16 | Reserved | | | | Sense Key | | | |
| 17 | Additional Sense Code | | | | | | | |
| 18 | Additional Sense Code Qualifier | | | | | | | |
| 19 | Vendor specific | | | | | | | |

Figure 76. Log Sense Page 10, Self-Test Results

- **Parameter Code** identifies the log parameter for the log page. The parameter Code field for the results of the most recent test will be 0001h. The parameter for the next most recent will be 0002h.
- **Function Code** contains the content of the Function Code field in the SEND DIAGNOSTIC command that initiated this self-test.
- **Self-Test Results Value** is described in the table below.

| Value | Description |
|---|---|
| 0h | The self-test routine completed without error. |
| 1h | The background self-test routine was aborted by the initiator using a SEND DIAGNOSTIC command with the Abort Background self-test function. |
| 2h | The self-test routine was aborted by the application client by a Task Management function of a reset. |
| 3h | An unknown error occurred while the target was executing the self-test routine and the target was unable to complete the self-test routine. |
| 4h | The self-test completed with a test element that failed and it is not known which test element failed. |
| 5h | The first segment of the self-test failed. |
| 6h | The second segment of the self-test failed. |
| 7h | The third or greater segment of the self-test failed (see the Extended segment number field). |
| 8h-Eh | Reserved. |
| Fh | The self-test is in progress. |

Figure 77. Log Sense Page 10 Self-Test Results

- **Extended Segment Number** This field will be used to identify the number of the segment that failed during self-test. If no segment failed, this field will be 00h.
- **Timestamp** This field contains the total accumulated power-on hours of the target at the time the self-test completed.
- **LBA of first failure** This field contains the LBA of the first logical block address where a self-test error occurred. If no errors occurred during the self-test or the error is not related to a Logical Block Address then the field will be FFFFFFFFFFFFFFFFh.
- **Sense Key, Additional Sense Code and Additional Sense Code Qualifier** These fields will contain the additional information relating to the error or exception conditions during self-test.

| Extended Segment Number | Short Self-Test | Extended Self-Test |
|---|---|---|
| 1h | Drive Ready Test | |
| 2h | RAM Test | |
| 3h | Drive Diagnostics | |
| 4h | SMART | |
| 5h | Low Level Format check | |
| 6h | Physical Head Check | |
| 7h | Error Log pages check | |
| 8h | Random Verify | |
| 9h | - Verify First 300 MB<br>- Verify Last 100 MB | Verify all LBA’s |
| Ah | Verify sectors adjacent to P & G list | |
| Bh | Recheck Error Log pages | |
| Ch | Recheck SMART | |

Figure 78. Log Sense Page 10 Extended Segment Number

See Section 10.17, "Diagnostics" on page 215 a for detailed listing of operations carried out by SEND DIAGNOSTIC command and Power on Diagnostics.

## 7.6.11 Log Sense Page 2F

This page contains SMART Status and Temperature Reading.

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | **7** | **6** | **5** | **4** | **3** | **2** | **1** | **0** |
| 0 | Reserved | | Page code = 2Fh | | | | | |
| 1 | Reserved | | | | | | | |
| 2-3 | PageLength = 8 | | | | | | | |
| 4-5 | Parameter Code = 0000h | | | | | | | |
| 6 | DU = 0 | DS = 0 | TSD = 0 | ETC = 0 | TMC = 0 | | LBIN = 1 | LP = 1 |
| 7 | Parameter Length = 04h | | | | | | | |
| 8 | SMART Sense Code Byte | | | | | | | |
| 9 | SMART Sense Qualifier | | | | | | | |
| 10 | Most Recent Temperature Reading | | | | | | | |
| 11 | Vendor HDA Temperature Trip Point | | | | | | | |

Figure 79. Log Sense Page 2F

# 7.7 MODE SENSE (1A)

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | **7** | **6** | **5** | **4** | **3** | **2** | **1** | **0** |
| 0 | Command Code = 1Ah | | | | | | | |
| 1 | Reserved | | | RSVD | DBD | Reserved = 0 | | |
| 2 | PCF | | Page Code | | | | | |
| 3 | Reserved = 0 | | | | | | | |
| 4 | Allocation Length | | | | | | | |
| 5 | VU = 0 | | Reserved = 0 | | | | FLAG | LINK |

Figure 80. MODE SENSE (1A)

The MODE SENSE (1A) command provides a means for the drive to report various device parameters to the initiator. It is the complement to the MODE SELECT command.

If the **DBD** (Disable Block Descriptor) bit is zero, the target will return the Block Descriptor. If the DBD bit is set to 1, the target will not return the Block Descriptor.

**Allocation Length** indicates the maximum number of bytes which the initiator has set aside for the DATA IN phase. A value of zero is not considered an error. If the allocation length is smaller than the amount available, that portion of the data up to the allocation length will be sent. This may result in only a portion of a multi-byte field being sent.

**Page Control Field:** PCF (Page Control Field) defines the type of Page Parameter values to be returned.

**PCF Meaning**

**0 0** **Report current values.** The drive returns the current values under which the logical unit is presently configured for the page code specified. The current values returned are

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

**Note:** The drive will not process the Mode Select command until the completion of spin-up. Therefore, the initiator cannot modify the current values prior to the saved values being read in.

**0 1** **Report changeable value.** The drive returns the changeable values for the page code specified. The page requested is returned containing information that indicates which fields are changeable. All bits of parameters that are changeable shall be set to one. Parameters that are *defined by the drive* shall be set to zero. If any part of a field is changeable, all bits in that field shall be set to one.

**Note:** For a value field such as the buffer ratios of page 2 the bit field will not indicate the range of supported values but rather that the field is supported.

**1 0** **Report default value.** The drive returns the default values for the page code specified. The parameters not supported by the drive are set to zero.

**1 1** **Report saved value.** The drive returns the saved value for the page code specified.

Saved values are one of following:

- the values saved as a result of MODE SELECT command
- identical to the default values
- zero when the parameters are not supported

The Page Length byte value of each page returned by the drive indicates up to which fields are supported on that page.

**Page Code:** This field specifies which page or pages to return. Page code usage is defined in the figure below.

| Page Code | Description |
|---|---|
| 00h - 1Ch<br>3Fh | Return specific page<br>Return all available pages |

Figure 81. Page Code Usage

## 7.7.1 Mode Parameter List

The mode parameter list contains a header followed by zero or more block descriptors followed by zero or more variable length pages.

### 7.7.1.1 Header

The six-byte command descriptor block header is defined below.

**Mode parameter header (6)**

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | **7** | **6** | **5** | **4** | **3** | **2** | **1** | **0** |
| 0 | Mode Data Length | | | | | | | |
| 1 | Medium Type = 0 | | | | | | | |
| 2 | WP | Reserved = 0 | | | | | | |
| 3 | Block Descriptor Length (= 0 or 8) | | | | | | | |

Figure 82. Mode parameter header (6)

The ten-byte command descriptor block header is defined below.

**Mode parameter header (10)**

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | **7** | **6** | **5** | **4** | **3** | **2** | **1** | **0** |
| 0<br>1 | (MSB) Mode Data Length (LSB) | | | | | | | |
| 2 | Medium Type = 0 | | | | | | | |
| 3 | WP | Reserved = 0 | | | | | | |
| 4<br>5 | Reserved = 0 | | | | | | | |
| 6<br>7 | (MSB) Block Descriptor Length ( = 0 or 8) (LSB) | | | | | | | |

Figure 83. Mode parameter header (10)

- **Mode Data Length**. When using the MODE SENSE command, the mode data length field specifies the length in bytes of the following data that is available to be transferred. The mode data length does not include the length byte itself. When using the MODE SELECT command, this field is reserved.
- **Medium Type** field is always set to zero in the drive (Default Medium Type).
- **WP**. When used with the MODE SELECT command, the Write Protect (WP) bit is reserved.

  When used with the MODE SENSE command, a Write Protect (WP) bit of zero indicates that the medium is write enabled.
- **Block Descriptor Length** specifies the length in bytes of the block descriptors.

  When used with the MODE SELECT command, zero or eight is supported by the drive.

  When used with the MODE SENSE command, the drive returns eight to indicate that only a single block descriptor is available.

### 7.7.1.2 Block Descriptor

| | |
|---|---|
| Byte 0<br>Byte 1<br>Byte 2<br>Byte 3 | Number of Blocks (MSB)<br><br><br>(LSB) |
| Byte 4 | Density code = 0 |
| Byte 5<br>Byte 6<br>Byte 7 | Block Length<br><br> |

Figure 84. MODE Parameter Block Descriptor

The Block descriptor provides formatting information about the Number of Blocks (user addressable) to format at the specified Block Length.

- **Number of Blocks**
  When used with the MODE SELECT command, the **Number of Blocks** field must be
  - Zero to indicate not to change available blocks
  - 0xFFFFFFFF to indicate all available blocks
  - The exact number of blocks in the data area of the drive, which can be obtained with the MODE SENSE
  - The number of blocks less than exact one, in order to **CLIP** the number of blocks

  Any other value is invalid and causes the command to fail with CHECK CONDITION status.

  When used with the MODE SENSE command, the field contain exact number of blocks.

- **Block Length**
  When used with the MODE SELECT command, the **Block length** field must contain the value from 512 to 528 (2 bytes step) or zero. Otherwise the drive will terminate the command with CHECK CONDITION status.

  A Format Unit command is required to cause these parameters to become current only if the block length parameter is different from the current block length.

  When used with the MODE SENSE command, the field is dependent on how the media is currently formatted.

### 7.7.1.3 Page Descriptor

| | | | |
|---|---|---|---|
| Byte 0 | PS | RSVD= 0 | Page Code |
| Byte 1 | Page Length | | |
| Byte 2-n | Mode Parameters | | |

Figure 85. MODE Parameter Page Format

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

When using the MODE SENSE command, a Parameter Savable (PS) bit of one indicates that the mode page can be saved by the drive in the reserved area of the drive. A PS bit of zero indicates that the supported parameters cannot be saved. When using the MODE SELECT command, the PS bit is reserved (zero).

The drive supports the following mode page code:

| Page | Description | PS |
|---|---|---|
| 00 | Vendor Unique Parameters | 1 |
| 01 | Read-Write Error Recovery Parameters | 1 |
| 02 | Disconnect/Reconnect Control Parameters | 1 |
| 03 | Format Device Parameters | 0 |
| 04 | Rigid Disk Geometry Parameters | 0 |
| 07 | Verify Error Recovery Parameters | 1 |
| 08 | Caching Parameters | 1 |
| 0A | Control Mode Page | 1 |
| 0C | Notch Parameters | 1 |
| 19 | Port Control Page | 1 |
| 1A | Power Control Parameters | 1 |
| 1C | Informational Exceptions Control | 1 |

Figure 86. Page Code Usage

The page length field specifies the length in bytes of the mode parameters that follow. If the initiator does not set this value to the value that is returned for the page by the MODE SENSE command, the drive will terminate the command with CHECK CONDITION status.

## 7.7.2 Mode Page 0 (Vendor Unique Parameters)

| Byte | Bit | | | | | | | | Default |
|---|---|---|---|---|---|---|---|---|---|
| | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | |
| 0 | PS | RSVD=0 | Page Code = 00h | | | | | | 80h |
| 1 | Page Length = 0Eh | | | | | | | | 0Eh |
| 2 | QPE | SSM | Ignored | UAI | MRG | Ignored | | ARHES | 11h |
| 3 | ASDPE | Ignored | CMDAC | Ignored | | | RRNDE | CPE | 21h |
| 4 | Ignored | | | | | | | | 00h |
| 5 | Ignored | | | FDD | Ignored | | CAEN | Ignored | 02h |
| 6 | IGRA | AVERP | Ignored | | | | | | 00h |
| 7 | Ignored | | | | | | | | 00h |
| 8 | Ignore | ADC | Ignored | | LED Mode | | | | 40h |
| 9 | Temperature | | | | | | | | 00h |
| 10 | Command Aging Limit (Hi byte) | | | | | | | | 00h |
| 11 | Command Aging Limit (Low byte) | | | | | | | | 30h |
| 12 | QPE Read Threshold | | | | | | | | 0Ah |
| 13 | QPE Write Threshold | | | | | | | | 0Ah |
| 14 | DRRT | Ignored | | | FFMT | Ignored | | | 00h |
| 15 | Ignored | | FCERT | Ignored | | | Reserved = 0 | | 00h |

Figure 87. Vendor Unique Parameters - Page 0

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

- **QPE** (Qualify Post Error) bit allows the Initiator to inhibit the reporting of recovered data errors which are recovered under Data Recovery Procedure (DRP) step. A QPE bit of zero causes the Target to report all recovered data errors. A QPE bit of one causes the Target to report only those recovered data errors which exceed the QPE threshold. (note: QPE = 1 is NOT supported.)
- **SSM** (Synchronous Select Mode) enables the drive to initiate SDTR and WDTR messages the first time the drive is selected with Attention. This is only allowed if the Enable TI-SDTR jumper is not installed. If the Enable TI-SDTR jumper is installed, this bit is ignored. (Refer to Section 6.2.1.5**,** "Position #9-10 **-** Enable TI-SDTR" on page 26.
- **UAI** (Unit Attention Inhibit) bit is not used during normal operation when the UAI jumper is removed from the drive. It may however be changed by the user with no effects. If the UAI jumper is added to the drive, then this bit controls the generation of unit attention conditions.
- **MRG** (Merge G-List into P-List) bit is set to 1 for merging the Grown Defect List (G-List) entries into the Primary Defect List (P-List) during Format Unit command.
- **ARHES** (Automatic Reassign Hard Error Sites) bit of one indicates that the drive will automatically reassign a hard read error. When the ARHES bit is active, the LBA with an unrecovered read error is internally registered as a reassign candidate and, when the write command is received to the failing

LBA, the data is written and verified if the error still exists. If the error still occurs on the failing LBA, it will be reassigned. If the error is cleared, the Target will remove the reassign candidate internally for the previously failing LBA. ARHES bit of zero indicates the drive will not automatically recover a hard read error. ARRE and AWRE bits in Mode Page 1 do not affect ARHES operation and ARHES works independently.

- **ASDPE** (Additional Save Data Pointer Enable) bit is used to control the sending of additional save data pointers messages. When set it will cause a save data pointers message to be sent on every disconnection. This bit is only used by the Target after the Default Mode parameter values are over-ridden with the Saved values which are read from the Reserved Area of the media as a part of the motor startup sequence. Before the Saved values are read from the Reserved Area of the media, the Save Data Pointer message is always sent to the Initiator prior to disconnection. When not set, a save data pointers message is sent only if the current connection contained a data phase and a further data phase will be required to complete the command.
- **CMDAC** (Command Active) bit in conjunction with **LED Mode** bits determines if an LED on the drive is activated while commands are active. If CMDAC bit is one and LED Mode bits are zero, an LED driver is active when a command is queued or executed.
- **RRNDE** (Report Recovered Non Data Errors) bit controls the reporting of recovered Non Data Errors when the PER bit is set. If RRNDE is set, recovered Non Data Errors are reported. If the RRNDE bit is not set, then recovered Non Data Errors are not reported.
- **CPE** (Concurrent Processing Enable) bit is allowed to be modified by the initiator for host system device driver compatibility. Read(6), Read extend(10), Write(6), Write extend(10), and untagged and unlinked Request Sense or Inquiry can be executed concurrently in both CPE bit set to 0 and 1.
- **FDD** (Format Degraded Disable) bit of 1 prevents the drive from reporting Format degraded. An FDD bit of 0 indicates that Format degraded is reported for the Test Unit Ready Command and causes media access commands (i.e. Read, Write) to report a media error if degraded.
- **CAEN** When set this bit causes the Command Age Limit timer to be used to avoid commands waiting in the command queue for an indefinite period. When commands have been in the queue for a period of time greater than the timer limit they will be reordered to be executed on a first come first served basis. When this bit is reset, commands are always executed based on the queue reordering rules.
- **IGRA** (Ignore Reassigned LBA) bit is set to 1 for preventing the drive from reassigned processing against reassigned LBA when RC bit (Mode Page 1 byte 2 bit 4) is active. The main purpose of this bit is to avoid undesirable read processing time delay due to reassigned LBA processing for continuous data availability requirements such as Audio Visual application. IGRA bit set to 0 specifies that the drive shall process reassigned LBA even if RC bit is active.
- **AVERP** (AV ERP Mode) bit is set to one in order to specify maximum retry counts during DRP and command execution time limit. When AVERP bit is set to one, the maximum retry counts for read and write operations are specified by Read Retry Count (Mode Page 1 Byte 3) and Write Retry Count (Mode Page 1 Byte 8) respectively. Recovery Time Limit (Mode Page 1 Byte 10 and 11) is effective to limit the command execution time. AVERP bit is set to zero to ignore the Recovery Time Limit value and to specify that the drive shall process DRP up to the default maximum retry count when Read Retry Count and Write Retry Count are set to a non-zero value.
- **ADC** (Adaptive Cache Enable), when set, allows the drive to modify the read-ahead caching algorithm, ignoring parameters in Page 8. The adaptation is based on analyzing the most recent command history and the current contents of the cache buffers.

- **LED Mode** is designed to control the operation of a drive LED driver.

  - LED Mode = 0000b

  The CMDAC bit controls the LED.
  CMDAC = 1 (Command Active)
  CMDAC = 0 (Motor Active)
  - LED Mode = 0001b (Motor Active)

  When the motor is spinning, the LED is high.
  - LED Mode = 0010b (Command Active).

  When there is a command active or in the queue, the LED is high.

- **Temperature Threshold** specifies the threshold value in degrees Celsius for the thermal sensor warning message. A value of 0 selects the default value (85 degrees Celsius).
- **Command Aging Limit** This value is used to control the maximum time a command should wait in the command queue when the CAEN bit is set. Each unit of this timer is 50 ms.
- **QPE Read Threshold** specifies the error reporting threshold for read operations when the QPE bit is set to one. (note: QPE = 1 is NOT supported.)
- **QPE Write Threshold** specifies the error reporting threshold for write operations when the QPE bit is set to one. (note: QPE = 1 is NOT supported.)
- **DRRT** (Disable Restore Reassign Target) bit disables the reading and restoration of the target LBA during a Reassign Blocks command. If the DRRT bit is zero, the reassign command attempts to restore the target LBA's data. If the data cannot be restored, the target LBA is reassigned and written with a data pattern of all 00's. If the DRRT bit is one, no attempt is made to restore the target LBA.
- **FFMT** (Fast Format enable) bit allows the formatting of the drive without any writes to the customer media. All format operations are allowed including changing block sizes and manipulating defects. The drive will operate normally after a fast format with the following caveat: since no data is written to any customer data blocks as a result of a Fast Format operation, there is a possibility that a read attempt to any particular block (without having previously written to that block) will result in an unrecoverable data error. This will most likely happen if the block size is changed as every LBA will contain data of an incorrect length and apparently an incorrect starting point. It is also possible to generate an uncorrectable data error without changing block sizes if the defect list is shortened and previously bad blocks become visible in the customer address space. Of course ALL DATA ON THE DRIVE WILL BE LOST as the result of any format operation and so any attempt to read blocks which have not been written to will result in unpredictable behavior.
- **FCERT** (Format Certification) bit determines whether the certification step will be performed during a Format Unit command. A FCERT bit set to 0 disables certification. A FCERT bit set to 1 enables the certification step.

## 7.7.3 Mode Page 1 (Read/Write Error Recovery Parameters)

| Byte | Bit | | | | | | | | Default |
|---|---|---|---|---|---|---|---|---|---|
| | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | |
| 0 | PS | RSVD=0 | Page Code = 01h | | | | | | 81h |
| 1 | Page Length = 0Ah | | | | | | | | 0Ah |
| 2 | AWRE | ARRE | TB | RC | EER=0 | PER | DTE | DCR | C0h |
| 3 | Read Retry Count | | | | | | | | 01h |
| 4 | Correction Span | | | | | | | | 00h |
| 5 | Head Offset Count (Ignored) | | | | | | | | 00h |
| 6 | Data Strobe Offset Count (Ignored) | | | | | | | | 00h |
| 7 | Reserved | | | | | | | | 00h |
| 8 | Write Retry Count | | | | | | | | 01h |
| 9 | Reserved | | | | | | | | 00h |
| 10 | (MSB) Recovery Time Limit | | | | | | | | 00h |
| 11 | (LSB) | | | | | | | | 00h |

Figure 88. Mode Page 1 (Vendor Unique Parameters )

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

- **AWRE**, an Automatic Write Reallocation Enabled bit, set to zero indicates that the drive shall not perform automatic reallocation of defective data blocks during write operations. An AWRE bit set to one indicates that the drive shall perform automatic reallocation of defective data blocks during write operations.
- **ARRE**, an Automatic Read Reallocation Enabled bit, set to zero indicates that the drive shall not perform automatic reallocation of defective data blocks during read operations. A ARRE bit set to one indicates that the drive shall perform automatic reallocation of defective data blocks during read operations.
- **TB**, a Transfer Block bit, set to one indicates that a data block that is not recovered within the recovery limits specified shall be transferred to the initiator before CHECK CONDITION status is returned. A TB bit set to zero indicates that such a data block shall not be transferred to the initiator. Data blocks that can be recovered within the recovery limits are always transferred regardless of the value of the bit.
- **RC**, a Read Continuous bit, set to one requests the Target to transfer the entire requested length of data without adding delays which would increase or ensure data integrity. This implies that the Target may send erroneous data. This bit has priority over all other error control bits (PER, DTE, DCR, TB). RC set to zero indicates normal interpretation of PER, DTE, DCR, and TB values. The RC bit setting is used by the Target when reporting errors associated with the transfer of the Initiator's data for the following commands:

  - Read(6)
  - Read(10)

For all other commands, the RC bit setting is unused and treated as zero.

**Note:** The Target implementation of the RC option is to disable error detection of the data fields but continue normal error detection and recovery for errors occurring in the servo field. If a servo field failure occurs, normal DRP could result in considerable recovery action, including proceeding through all levels of DRP.

- **EER**, an Enable Early Recovery bit, **must be set to zero**, indicating that the drive shall use an error recovery procedure that minimizes the risk of misdetection or miscorrection during the data transfer. Data shall not be fabricated.
- **PER**, a Post Error bit, is set to one to indicate that the drive reports recovered errors.
- **DTE**, a Disable Transfer on Error bit, must be set to 0 to indicate that the drive continues the DATA phase upon detection of a recovered error .
- **DCR**, a Disable Correction bit, is set to one to indicate that Error Correction Code is not used for data error recovery. A DCR bit of zero indicates that ECC is applied to recover the data.
- **Read Retry Count** sets a limit on the amount of DRP passes in which the Target attempts to recover read errors. A value of zero disables all data recovery procedures. When AVERP bit (Mode Page 0 Byte 6 Bit 6) is zero, a value of non-zero in Read Retry Count enables all steps of DRP. When AVERP bit is one, the number in Read Retry Count sets the maximum retry count of DRP.
- **Correction Span** field specifies the size in bits of the largest data error burst for which data error correction may be attempted. Any value can be placed in this field including zero. The drive will always use its default correction capabilities.
- **Head Offset Count** is not supported by the drive.

  **Note:** Head Offset is implemented in the read error recovery routine. The user can not modify the offset value.
- **Write Retry Count** sets a limit on the amount of DRP passes in which the Target attempts to recover write errors. A value of zero disables all data recovery procedures. When AVERP bit (Mode Page 0 Byte 6 Bit 6) is zero, a value of non-zero in Write Retry Count enables all steps of DRP. When AVERP bit is one, the number in Write Retry Count sets the maximum retry count of DRP.
- **Recovery Time Limit** indicates the period in 1 millisecond increments for the maximum command execution time.The value must be from 40 ms to 65535 ms (65.5 seconds). If a command is not able to be completed within the limit, a check condition will be returned when following mode page parameters are set.

  - WCE = 0

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

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

**1 0 0 0** The highest level error is reported at the end of transfer. Retries and error correction are attempted. Recovered or corrected data (if any) or both are transferred with CHECK CONDITION status and RECOVERED ERROR Sense Key set at the end of the transfer.

**no err** The transfer length is exhausted.

**soft err** The transfer length is exhausted. Transferred data includes blocks containing recovered errors. The information byte in the sense data will contain the logical block address of the last recovered error.

**hard err** Data transfer stops when an unrecoverable error is encountered. The unrecoverable block is not transferred to the initiator. The drive then creates the CHECK CONDITION status with the appropriate Sense Key.

**1 0 0 1** The highest level error is reported at the end of transfer. Retries and error correction are attempted. Recovered or corrected data (if any) or both are transferred with CHECK CONDITION status and RECOVERED ERROR Sense Key set at the end of the transfer.

**no err** The transfer length is exhausted.

**soft err** The transfer length is exhausted. Transferred data includes blocks containing recovered errors. The information byte in the sense data will contain the logical block address of the last recovered error.

**hard err** Data transfer stops when an unrecoverable error is encountered. The unrecoverable block is transferred to the initiator. The drive then creates the CHECK CONDITION status with the appropriate Sense Key.

**1 0 1 0** The highest level error is reported at the end of transfer. Retries are attempted but ECC is not applied. Recovered or corrected data (if any) or both are transferred with CHECK CONDITION status and RECOVERED ERROR Sense Key set at the end of the transfer.

**no err** The transfer length is exhausted.

**soft err** The transfer length is exhausted. Transferred data includes blocks containing recovered errors. The information byte in the sense data will contain the logical block address of the last recovered error.

**hard err** Data transfer stops when an unrecoverable error is encountered. The unrecoverable block is not transferred to the initiator. The drive then creates the CHECK CONDITION status with the appropriate Sense Key.

**1 0 1 1** The highest level error is reported at the end of transfer. Retries and error correction are attempted. Recovered or corrected data (if any) or both are transferred with CHECK CONDITION status and RECOVERED ERROR Sense Key set at the end of the transfer.

**no err** The transfer length is exhausted.

**soft err** The transfer stops on the first soft error detected. The information in the sense data shall contain the LBA of the block in error.

**hard err** Data transfer stops on the unrecoverable error. The unrecoverable error block is not returned to the initiator. The drive then creates the CHECK CONDITION status with the appropriate Sense Key.

**1 1 0 0** The highest level error is reported at the end of transfer. Retries and error correction are attempted. Recovered or corrected data (if any) or both are transferred with CHECK CONDITION status and RECOVERED ERROR Sense Key set at the end of the transfer.

**no err** The transfer length is exhausted.

**soft err** The transfer stops on the first soft error detected. The information in the sense data shall contain the LBA of the block in error.

**hard err** Data transfer stops on the unrecoverable error. The unrecoverable error block is not returned to the initiator. The drive then creates the CHECK CONDITION status with the appropriate Sense Key.

**1 1 0 1** The highest level error is reported at the end of transfer. Retries and error correction are attempted. Recovered or corrected data (if any) or both are transferred with CHECK CONDITION status and RECOVERED ERROR Sense Key set at the end of the transfer.

**no err** The transfer length is exhausted.

**soft err** The transfer stops on the first soft error detected. The information in the sense data shall contain the LBA of the block in error.

**hard err** Data transfer stops on the unrecoverable error. The unrecoverable error block is returned to the initiator. The drive then creates the CHECK CONDITION status with the appropriate Sense Key.

**1 1 1 0** The highest level error is reported at the end of transfer. Retries are attempted but ECC is not applied. Recovered data are transferred with CHECK CONDITION status and RECOVERED ERROR Sense Key set at the end of the transfer.

**no err** The transfer length is exhausted.

**soft err** The transfer stops on the first soft error detected. The recovered error block is returned to the initiator. The information in the sense data shall contain the logical block address of the block in error.

**hard err** Data transfer stops on the unrecoverable error. The unrecoverable error block is not returned to the initiator The drive then creates the CHECK CONDITION status with the appropriate Sense Key.

**1 1 1 1** The highest level error is reported at the end of transfer. Retries are attempted but ECC in not applied. Recovered or corrected data (if any) or both are transferred with CHECK CONDITION status and RECOVERED ERROR Sense Key set at the end of the transfer.

**no err** The transfer length is exhausted.

**soft err** The transfer stops on the first soft error detected. The information in the sense data shall contain the logical block address of the block in error.

**hard err** Data transfer stops on the unrecoverable error. The unrecoverable error block is returned to the initiator. The drive then creates the CHECK CONDITION status with the appropriate Sense Key.

## 7.7.4 Mode Page 2 (Disconnect/Reconnect Parameters)

| Byte | Bit | | | | | | | | Default |
|---|---|---|---|---|---|---|---|---|---|
| | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | |
| 0 | PS | RSVD=0 | Page Code = 02h | | | | | | 82h |
| 1 | Page Length = 0Eh | | | | | | | | 0Eh |
| 2 | Read Buffer Full Ratio | | | | | | | | 00h |
| 3 | Write Buffer Empty Ratio | | | | | | | | 00h |
| 4-5 | (MSB) Bus Inactivity Limit = 0 (LSB) | | | | | | | | 00h<br>00h |
| 6-7 | (MSB) Disconnect Time Limit = 0 (LSB) | | | | | | | | 00h<br>00h |
| 8-9 | (MSB) Connect Time Limit = 0 (LSB) | | | | | | | | 00h<br>00h |
| 10-11 | (MSB) Maximum Burst Size (LSB) | | | | | | | | 00h<br>00h |
| 12 | RSVD | Fair arbitration | | | DIMM | RSVD | DTDC | | 00h |
| 13-15 | Reserved = 0 | | | | | | | | 00h<br>00h<br>00h |

Figure 89. Mode Page 2 (Disconnect/Reconnect Parameters )

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

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

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

- **Read Buffer Full Ratio** is the numerator of a fraction whose denominator is 256. The fraction indicates how full the drive data buffer should be before attempting to reconnect to the SCSI bus. If the ratio is set to 0h, the target will calculate and use an optimal ratio based on the negotiated transfer rate.
- **Write Buffer Empty Ratio** is the numerator of a fraction whose denominator is 256. The fraction indicates how empty the drive data buffer should be before attempting to reconnect to the SCSI bus. If the ratio is set to 0h, the target will calculate and use an optimal ratio based on the negotiated transfer rate.
- **Maximum Burst Size** is the maximum amount of data that the Target transfers during a data phase before disconnecting if the Initiator has granted the disconnect privilege. This value is expressed in increments of single block size (for example, a value of 0001h means 512 bytes, 0002h means 1024 bytes when the block size is 512 bytes). Disconnections attempted by the Target are on block boundaries only. For the case when (Maximum Burst Size x Block Size) is less than the Block Length, the Target will transfer 1 block of data before attempting to disconnect. Value of 0000h indicates there is

no limit on the amount of data transferred per connection. Regardless of the value in Maximum Burst Size the Target disconnects prior to completion of the data phase if the internal data buffer segment becomes empty during a Read command or full during a Write command.

- **DIMM** (Disconnect Immediate) bit allows the Initiator to determine whether the drive is required to disconnect after the receipt of a command and prior to starting a data phase. A DIMM bit of one indicates that the drive is required to disconnect after receiving a command, prior to starting a data phase. A DIMM bit of zero indicates that the drive may transfer data for a command immediately after receiving it, without disconnecting. Whether or not the drive does so depends upon the disconnect priviledges extended by the initiator, the workload seen at the drive, and the settings of the other parameters in this mode page.

  **Note:** Priority commands do not disconnect from the SCSI bus.

- **DTDC** (Data Transfer Disconnect Control) field defines further restrictions for when a disconnect is permitted.

  - A value of 00b indicates that DTDC is not used by the Target and the disconnect is controlled by the other fields in this page.
  - A value of 01b indicates that the target shall not attempt to disconnect once the data transfer of a command has started until all data which the command is to transfer has been transferred. The connect time limit and bus inactivity limit are ignored during the data transfer.
  - The value 10b is reserved.
  - A value of 11b indicates that the target shall not attempt to disconnect once the data transfer of a command has started until the command is complete. The connect time limit and bus inactivity limit are ignored once data transfer has started.

  **Note:** If DTDC is nonzero and the maximum burst size is nonzero, a CHECK CONDITION status will be returned. The sense key shall be set to ILLEGAL REQUEST and the additional sense code set to ILLEGAL FIELD IN PARAMETER LIST.

- **Fair arbitration** If the Fair arbitration filed is set to 000b, the target shall not use arbitration fairness during normal arbitration. If this field is set to a nonzero value, the target shall use arbitration fairness during normal arbitration.

Both the Read Buffer Full Ratio and the Write Buffer Empty Ratio pertain to the current active notch. For each active notch as defined in page 0Ch there are separate Read Buffer Full Ratios and Write Buffer Empty Ratios. When the active notch is set to zero, the values are applied in mode page 0Ch across all notches.

## 7.7.5 Mode Page 3 (Format Device Parameters)

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

Figure 90. Mode Page 3 (Format Device Parameters)

The format device page contains parameters which specify the medium format. This page contains no changeable parameters.

**Tracks per Zone** specifies the number of tracks within the zone. This field is a function of the active notch.

A value of 0 in the following parameters indicates that they are Target specific.

- Alternate Sectors per Zone
- Alternate Tracks per Zone
- Alternate Tracks per Logical Unit

**Sectors per Track** specifies the number of physical sectors within each track. This field is a function of the active notch.

**Data Bytes per Physical Sector** specifies the number of user data bytes per physical sector. The value depends upon the current formatted Block Length.

**Track Skew Factor** indicates the number of physical sectors between the last block of one track and the first block on the next sequential track of the same cylinder. This field is a function of the active notch.

**Cylinder Skew Factor** indicates the number of physical sectors between the last block of one cylinder and the first block on the next sequential cylinder. This field is a function of the active notch.

**SSEC** = Zero indicates that the drive does not support soft sector formatting.

**HSEC** = One indicates that the drive supports hard sector formatting.

**RMB** = Zero indicates that the media does not support removable. Fixed Disk.

**SURF** = Zero indicates that progressive addresses are assigned to all logical blocks in a cylinder prior to allocating addresses within the next cylinder.

## 7.7.6 Mode Page 4 (Rigid Disk Drive Geometry Parameters)

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

Figure 91. Mode Page 4 (Rigid Disk Drive Geometry Parameters**)**

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

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

## 7.7.7 Mode Page 7 (Verify Error Recovery Parameters)

| Byte | Bit | | | | | | | | Default |
|---|---|---|---|---|---|---|---|---|---|
| | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | |
| 0 | PS | RSVD=0 | Page Code = 07h | | | | | | 87h |
| 1 | Page Length = 0Ah | | | | | | | | 0Ah |
| 2 | Reserved = 0 | | | | EER=0 | PER | DTE=0 | DCR | 00h |
| 3 | Verify Retry Count | | | | | | | | 01h |
| 4 | Correction Span = 00h | | | | | | | | 00h |
| 5 | Reserved = 0 | | | | | | | | 00h |
| 6 | Reserved = 0 | | | | | | | | 00h |
| 7 | Reserved = 0 | | | | | | | | 00h |
| 8 | Reserved = 0 | | | | | | | | 00h |
| 9 | Reserved = 0 | | | | | | | | 00h |
| 10 | (MSB) Verify Recovery Time Limit (Not Used) | | | | | | | | 00h |
| 11 | (LSB) | | | | | | | | 00h |

Figure 92. Mode Page 7 (Verify Error Recovery Parameters)

The Verify recovery parameters are used by the Target when recovering from and reporting errors associated with the verification of the initiator's Data for the following commands:

- Verify
- Write and Verify - the verify portion of the command only.

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

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

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

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

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

## 7.7.8 Mode Page 8 (Caching Parameters)

| Byte | Bit | | | | | | | | Default |
|---|---|---|---|---|---|---|---|---|---|
| | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | |
| 0 | PS | RSVD=0 | Page Code = 08h | | | | | | 88h |
| 1 | Page Length = 12h | | | | | | | | 12h |
| 2 | IC | ABPF | CAP | DISC | SIZE | WCE | MF | RCD | 04h |
| 3 | Demand Read Retention Priority | | | | Write Retention Priority | | | | 00h |
| 4-5 | Disable Pre-fetch Transfer Length | | | | | | | | FFh<br>FFh |
| 6-7 | Minimum Pre-fetch | | | | | | | | 00h<br>00h |
| 8-9 | Maximum Pre-fetch | | | | | | | | FFh<br>FFh |
| 10-11 | Maximum Pre-fetch Ceiling | | | | | | | | FFh<br>FFh |
| 12 | FSW | LBCSS | DRA | Reserved = 0 | | | | | 00h |
| 13 | Number of Cache Segments | | | | | | | | 1Bh |
| 14-15 | (MSB) Cache Segment Size (LSB) | | | | | | | | 00h<br>00h |
| 16 | Reserved = 0 | | | | | | | | 00h |
| 17<br>18,19 | (MSB) Non Cache Segment Size (LSB) | | | | | | | | 000000h |

Figure 93. Page 8 (Caching Parameters)

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

- **IC** (Initiator Control) is internally ignored by the Target. The Target will always use the Number of Cache Segments or Cache Segment Size fields.
- **ABPF** (Abort Pre-Fetch) is not supported. The Target aborts the pre-fetch based on the internal algorithm.
- **CAP** (Caching Analysis Permitted) is not supported and is internally ignored by the Target.
- **DISC** (Discontinuity) is not supported and is internally ignored by the Target.
- **SIZE** (Size Enable) bit when set to one indicates that the Cache Segment Size is to be used to control caching segmentation. When SIZE is set to zero, the initiator requests that the Number of Cache Segments is to be used to control caching segmentation.
- **WCE** (Write Cache Enable) bit when set at zero indicates that the drive must issue Good Status for Write(6) or Write extend(10) command only after successfully writing the data to the media. When the WCE bit is set to one, the drive may issue Good Status for a Write(6) or Write extend(10)

command after successfully receiving the data but before writing it to the media. When WCE = 1, the drive operates as if AWRE = 1.

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

- **MF** (Multiplication Factor) determines how the Maximum Pre-fetch field is interpreted. This bit is ignored.
- **RCD** (Read Cache Disable) bit set at zero indicates that the drive may return some or all of the data requested by a Read (6) or Read (10) command by accessing the data buffer, not the media. An RCD bit set at one indicates that the Target does not return any of the data requested by a Read (6) or Read (10) command by accessing the data buffer. All of the data requested is read from the media instead.
- **Demand Read Retention Priority** sets the Retention Priority of data requested on a Read Command. This field is ignored.
- **Write Retention Priority** sets the Retention Priority of data provided on a Write Command. This field is ignored.
- **Disable Pre-fetch Transfer Length** specifies a number of LBA's which if exceeded by a read command length will cause the drive not to perform read ahead buffering after the command has completed. This field is ignored.
- **Minimum Pre-fetch** specifies the minimum number of LBA's that the drive should read ahead after each read command. A value of zero indicates that read ahead should be terminated immediately after a new command arrives, except when the new command is on the current head and track.
- **Maximum Pre-fetch** specifies the maximum number of LBA's to read ahead after a read command. This field is ignored.
- **Maximum Pre-fetch ceiling** specifies the maximum number of blocks the drive should attempt to read ahead. This field is ignored.
- **FSW** (Force Sequential Write) is not supported and internally ignored by the Target.
- **LBCSS** (Logical Block Cache Segment Size) bit when set to one indicates that the Cache Segment Size field units shall be interpreted as logical blocks. When it is set to zero, the Cache Segment Size field units shall be interpreted as bytes.
- **DRA** (Disable Read Ahead) bit when set to one requests that the target not read into the buffer any logical block beyond the addressed logical block(s). When it is set at zero, the target may continue to read logical blocks into the buffer beyond the addressed logical block(s).
- **Number of Cache Segments** field is used to select the number of data buffer cache segments. This parameter is valid only when the SIZE bit is set at zero. It is ignored when SIZE is set at one.

  The target supports the following Cache Segment configuration:

| Number of Segments | Segment Size (bytes) |
|---|---|
| 6 | 512K |
| 13 | 256K |
| 27 | 128K |

- **Cache Segment Size** field indicates the requested segment size in Bytes or Blocks, depending upon the value of the LBCSS bit. The Cache Segment Size field is valid only when the SIZE bit is one. It is ignored when SIZE is set at zero. The drive considers this a minimum size. It will be rounded to a value of 6, 13 or 27 segments.

## 7.7.9 Mode Page A (Control Mode Page Parameters)

| Byte | Bit | | | | | | | | Default |
|---|---|---|---|---|---|---|---|---|---|
| | **7** | **6** | **5** | **4** | **3** | **2** | **1** | **0** | |
| 0 | PS | RSVD=0 | Page Code = 0Ah | | | | | | 8Ah |
| 1 | Page Length = 0Ah | | | | | | | | 0Ah |
| 2 | Reserved = 0 | | | | | | | RLEC=0 | 00h |
| 3 | Queue Algorithm Modifier | | | | Rsvd=0 | QErr | | DQue | 00h |
| 4 | EECA=0 | Reserved = 0 | | | | RAENP = 0 | UAAENP = 0 | EAENP = 0 | 00h |
| 5 | Reserved = 0 | | | | | | | | 00h |
| 6-7 | (MSB) Ready AEN Holdoff Period (LSB) | | | | | | | | 00h |
| 8-9 | (MSB) Busy Timeout Period (LSB) | | | | | | | | 00h |
| 10-11 | (MSB) Extended Self-test Routine Completion Time (LSB) | | | | | | | | --- |

Figure 94. Page A (Control Mode Page Parameters)

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

- **Queue algorithm modifier** specifies restrictions on the algorithm used for reordering commands that are tagged with the SIMPLE message.

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

  1h : Unrestricted reordering allowed. The target may reorder the actual execution sequence of the queued commands in any manner it selects. Any data integrity exposures related to command sequence order are explicitly handled by the initiator through the selection of appropriate commands and queue tag messages.

  2h-7h : RESERVED.

  8 : Command reordering is disabled

  9-Fh : RESERVED

- **QErr** (Queue Error Management) The queue error management (QERR) field specifies how the device server shall handle blocked tasks when another task receives a CHECK CONDITION status.

| QERR value | Description |
|---|---|
| 00b | Specifies that all tasks from all initiators are blocked from execution when a Contintent Alligience (CA condition) is pending. Those blocked tasks are allowed to resume execution in a normal fashion after the CA condition is cleared. |
| 01b | Specifies that all tasks from all initiators are aborted when the Target returns Check Condition Status. A unit attention condition will be generated for each initiator which had commands in the queue except for the initiator that received the Check Condition Status. The sense key will be set to Unit Attention and the additional sense code will be set to COMMANDS CLEARED BY ANOTHER INITIATOR. |
| 10b | Reserved |
| 11b | Blocked tasks in the task set belonging to the initiator to which a CHECK CONDITION status is sent shall be aborted when the status is sent. |

- **DQue** (Disable Queuing) bit set at zero specifies that tagged queuing shall be enabled if the target supports tagged queuing. A DQue bit set at one specifies that tagged queuing shall be disabled. Any queue commands for that I_T_L nexus shall be aborted. Any subsequent queue tag message received shall be rejected with a MESSAGE REJECT message and I/O process shall be executed as an untagged command.
- **Ready AEN Holdoff Period** is ignored internally by the Target.
- **Busy Timeout Period** is ignored internally by the Target.
- **Extented Self-test Routine Completion Time** is an advisory parameter that an initiator may use to determine the time in seconds that the target requires to complete self-test routine when the target is not interrupted by an initiator and no errors occur during execution of the self-test routine.

## 7.7.10 Mode Page 0C (Notch Parameters)

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | **7** | **6** | **5** | **4** | **3** | **2** | **1** | **0** |
| 0 | PS = 1 | RSVD= 0 | Page Code = 0Ch | | | | | |
| 1 | Page Length = 16h | | | | | | | |
| 2 | ND = 1 | LPN = 0 | RSVD = 0 | | | | | |
| 3 | Reserved = 0 | | | | | | | |
| 4 | (MSB) Maximum Number of Notches = 0Bh | | | | | | | |
| 5 | (LSB) | | | | | | | |
| 6 | (MSB) Active Notch | | | | | | | |
| 7 | (LSB) | | | | | | | |
| 8 \| 11 | (MSB) Starting Boundary (LSB) | | | | | | | |
| 12 \| 15 | (MSB) Ending Boundary (LSB) | | | | | | | |
| 16 \| 23 | (MSB) Pages Notched = 0000000000000100Ch (LSB) | | | | | | | |

Figure 95. Page 0C (Notch Parameters)

The notch page contains parameters for direct-access devices which implement a variable number of blocks per cylinder. Each section of the logical unit with a different number of blocks per cylinder is referred as a notch. The only field that is changeable is the **Active Notch** field.

**ND** = ONE meaning that this device is a notched drive.

**LPN** = ZERO meaning that the notches are based upon physical parameters of the drive (cylinder #), not logical parameters.

**Maximum Number of Notches** is the number of notches the drive can support.

**Active Notch** indicates to which notch subsequent Mode Select/Sense command parameters pertain. A value of 0 is used for parameter values which apply to all notches. Values from 1 to the maximum value depending on the model specify the notch number, where notch 1 is the outermost notch. Following mode parameters are based on the current active notch:

- Mode Page 2
  - Read Buffer Full Ratio
  - Write Buffer Empty Ratio
- Mode Page 3
  - Alternate Sector per Zone

- Alternate Track per Zone
- Alternate Track per Logical Unit
- Sector per Track
- Track Skew Factor
- Cylinder Skew Factor

**Starting Boundary** contains the first physical location of the active notch. The first three bytes are the cylinder number and the last byte is the head. The value sent in this field is ignored.

**Ending Boundary** contains the last physical location of the active notch. The first three bytes are the cylinder number and the last byte is the head. The value sent in this field is ignored.

**Pages Notched** 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.

## 7.7.11 Mode Page 19 (Port Control)

| Byte | Bit | | | | | | | | Default |
|---|---|---|---|---|---|---|---|---|---|
| | **7** | **6** | **5** | **4** | **3** | **2** | **1** | **0** | |
| 0 | PS | RSVD | Page Code = 019h | | | | | | 99h |
| 1 | Page Length = 06h | | | | | | | | 06h |
| 2 | Reserved = 0 | | | | | | | | 00h |
| 3 | Reserved = 0 | | | | Protocol identifier = 1 | | | | 01h |
| 4<br>\|<br>5 | (MSB)<br>Synchronous transfer timeout<br>(LSB) | | | | | | | | 00h<br>\|<br>00h |
| 6 | Reserved = 0 | | | | | | | | 00h |
| 7 | Reserved = 0 | | | | | | | | 00h |

Figure 96. Page 19 (Port Control)

- **Protocol identifier** filed has a value of 1h to indicate SPI SCSI devices.
- **Synchronous transfer timeout** field must be 0000h.

## 7.7.12 Mode Page 1A (Power Control)

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | **7** | **6** | **5** | **4** | **3** | **2** | **1** | **0** |
| 0 | RSVD = 0 | RSVD = 0 | Page Code = 1Ah | | | | | |
| 1 | Page Length = 0Ah | | | | | | | |
| 2 | Reserved = 00h | | | | | | | |
| 3 | Reserved = 0 | | | | | | Idle | Standby |
| 4 \| 7 | (MSB) Idle Condition Timer = 0 (LSB) | | | | | | | |
| 8 \| 11 | (MSB) Standby Condition Timer (LSB) | | | | | | | |

Figure 97. Page 1A (Power Control)

- The value for the **Idle Condition Timer** and the corresponding **Idle** control bit is accepted. Despite this, the drive does not alter the internal power saving algorithms based upon this value.
- **Standby** bit of one indicates that the target shall use the **Standby Condition Timer** to determine the length of inactivity time to wait before entering the Standby condition. A standby bit of zero indicates that the target shall not enter the Standby condition.
- **Standby Condition Timer** field indicates the inactivity time in 100 millisecond increments that the target shall wait before entering the Standby condition. The minimum allowable inactivity time is 60 minutes. Any value less than this is accepted, but will automatically default to 60 minutes.

## 7.7.13 Mode Page 1C (Informational Exceptions Control)

| Byte | BIT | | | | | | | | Default |
|---|---|---|---|---|---|---|---|---|---|
| | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 | |
| 0 | PS | RSVD=0 | Page Code = 1Ch | | | | | | 9Ch |
| 1 | Page Length = 0Ah | | | | | | | | 0Ah |
| 2 | PERF | RSVD | EBF | EWASC | DEXCPT | TEST | RSVD | LOGERR | 00h |
| 3 | Reserved = 0 | | | | Method of Reporting | | | | 00h |
| 4 | (MSB) Interval Timer | | | | | | | | 00h |
| 5 | | | | | | | | | 00h |
| 6 | | | | | | | | | 00h |
| 7 | (LSB) | | | | | | | | 00h |
| 8 | (MSB) Report Count | | | | | | | | 00h |
| 9 | | | | | | | | | 00h |
| 10 | | | | | | | | | 00h |
| 11 | (LSB) | | | | | | | | 00h |

Figure 98. Page 1C (Informational Exceptions Control)

- **LOGERR** (Log Errors) is not used and ignored internally by the Target.
- **EWASC** (Enable Warning ASC) bit of zero indicates that Temperature Warning will not be reported. An EWASC bit of one allows Temperature Warning to be reported when DEXCPT bit is set to zero. If the temperature reading inside of DE exceeds the threshold (default or set by Mode Select Page 0), Recovered Error (Sense Key 01h, ASC 0Bh, ASCQ 01h) will be reported.

  DEXCPT bit of zero indicates that informational exception operations shall be enabled. The reporting of informational exception conditions when the DEXCPT bit is set to zero is determined from the method of reporting informational exceptions field.

  **Note:** Disable exception control (DEXCPT) bit of zero indicates the failure prediction threshold exceeded reporting shall be enabled. The method for 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.
- **DEXCPT** (Disable Exception Control) bit of zero indicates information exception operations are enabled. The reporting of information exception conditions when the DEXCPT bit is set to zero is determined from the Method of Reporting field. A DEXCPT bit of one indicates the Target disabled all information exception operations.
- **TEST** bit of one instructs the drive to generate false drive notifications at the next interval time, (as determined by the INTERVAL TIMER field), if the DEXCPT is zero. The Method of Reporting and Report Count would apply. The false drive failure is reported as sense qualifier 5DFFh. The TEST bit of zero instructs the drive to stop generating any false drive notifications.

- **EBF** (Enable Background Function) bit of one indicates that the target enables background functions. An EBF bit of zero indicates that the target disables the functions.
- **PERF** (Performance) bit is not used and ignored internally by the Target.
- **Method of Reporting** Informational Exceptions indicates the methods used by the Target to report informational exception conditions.

| Code | Description |
|---|---|
| **0h** | **No reporting of informational exception condition:** This method instructs the target to not report informational exception condition. |
| **1h** | Asynchronous event reporting: Not supported. |
| **2h** | **Generate unit attention:** This method instructs the target to report informational exception conditions by returning a CHECK CONDITION status on any command. The sense key is set to UNIT ATTENTION and the additional sense code indicates the cause of the informational exception condition. The command that has the CHECK CONDITION is not executed before the informational exception condition is reported. |
| **3h** | **Conditionally generate recovered error:** This method instructs the target to report informational exception conditions, dependent on the value of the PER bit of the error recovery parameters mode page, by returning a CHECK CONDITION status on any command. The sense key is set to RECOVERED ERROR and the additional sense code indicates the cause of the informational exception condition. The command that has the CHECK CONDITION completes without error before any informational exception condition is reported. |
| **4h** | **Unconditionally generate recovered error:** This method instructs the target to report informational exception conditions, regardless of the value of the PER bit of the error recovery parameters mode page, by returning a CHECK CONDITION status on any command. The sense key is set to RECOVERED ERROR and the additional sense code indicates the cause of the informational exception condition. The command that has the CHECK CONDITION completes without error before any informational exception condition is reported. |
| **5h** | **Generate no sense:** This method instructs the target to report informational exception conditions by returning a CHECK CONDITION status on any command. The sense key is set to NO SENSE and the additional sense code indicates the cause of the informational exception condition. The command that has the CHECK CONDITION completes without error before any informational exception condition is reported. |
| **6h** | **Only report informational exception condition on request:** This method instructs the target to preserve the informational exception(s) information. To find out about information exception conditions the Application Client polls the target by issuing an unsolicited REQUEST SENSE command. The sense key is set to NO SENSE and the additional sense code indicates the cause of the informational exception condition. |
| **7h-Fh** | Reserved. |

- **Interval Timer** field indicates the period in 100 millisecond increments for reporting that an informational exception condition has occurred. The target shall not report informational exception conditions more frequently than the time specified by the Interval Timer field and as soon as possible after the time interval has elapsed. After the informational exception condition has been reported the interval timer is restarted. A value of zero in the Interval Timer field indicated that the target only reports the informational exception condition one time.
- **Report Count** field indicates the number of times the target reports an informational exception condition. The Report Count of ZERO indicates no limits on the number of times the target reports an informational exception condition.

# 7.8 MODE SENSE (5A)

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
| 0 | Command Code = 5Ah | | | | | | | |
| 1 | Reserved = 0 | | | RSVD | DBD | Reserved = 0 | | |
| 2 | PCF | | Page Code | | | | | |
| 3 | Reserved = 0 | | | | | | | |
| 4 | Reserved = 0 | | | | | | | |
| 5 | Reserved = 0 | | | | | | | |
| 6 | Reserved = 0 | | | | | | | |
| 7-8 | (MSB) Allocation Length (LSB) | | | | | | | |
| 9 | VU = 0 | | Reserved = 0 | | | | FLAG | LINK |

Figure 99. MODE SENSE (5A)

The MODE SENSE (5A) command provides a means for the drive to report various device parameters to the initiator. See the MODE SENSE (1A) command for a description of the fields in this command.

# 7.9 MODE SELECT (15)

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | **7** | **6** | **5** | **4** | **3** | **2** | **1** | **0** |
| 0 | Command Code = 15h | | | | | | | |
| 1 | Reserved = 0 | | | PF=1 | Reserved = 0 | | | SP |
| 2<br>3 | Reserved = 0 | | | | | | | |
| 4 | Parameter List Length | | | | | | | |
| 9 | VU = 0 | | Reserved = 0 | | | | FLAG | LINK |

Figure 100. MODE SELECT (15)

The MODE SELECT (15) command provides a means for the initiator to specify LUN or device parameters to the Target. It also allows an Initiator to specify options the Target uses in error recovery and Caching.

There is a single set of Mode Page parameters shared by all initiators.

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

**SP** Pages. This indicates

0 The drive shall not save the pages sent during the Data Out phase but will use them for all following commands until the power is removed, a reset is received, or a new mode select command is received.

1 The drive will save the data in the reserved area of the disk. It will be used for all the following commands until another mode select command is issued. This information is maintained over a power cycle or reset of the drive.

**Parameter List Length**

This specifies the number of bytes to be sent from the initiator. A parameter list length of zero suppresses data transfer and is not considered as an error.

The MODE SELECT parameter list contains a four-byte header, followed by zero or one block descriptor followed by zero or more pages. The pages which are valid with this command are defined in the addendum under the heading **Mode Select Data** as they vary with the drive model.

### Application Note

The initiator should issue a MODE SENSE command requesting all Changeable values (see PCF field in byte two of the CDB in **MODE SENSE (1A)** on page 82) prior to issuing a MODE SELECT command. This is necessary to find out which pages are implemented by the drive and the length of those pages. In the Pages of the MODE SENSE command the drive will return the number of bytes supported for each Page. The Page Length set by the initiator in the MODE SELECT command must be same value as returned by the drive in MODE SENSE Page Length. If not, the drive will return CHECK CONDITION status with sense key of ILLEGAL REQUEST.

**Note:** If an initiator sends a MODE SELECT command that changes any parameters that apply to other initiators, the drive shall generate an unit attention condition for all initiators except for the one that issued the MODE SELECT command. The drive shall set the additional sense code to PARAMETERS CHANGED (2Ah).

# 7.10 MODE SELECT (55)

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | **7** | **6** | **5** | **4** | **3** | **2** | **1** | **0** |
| 0 | Command Code = 55h | | | | | | | |
| 1 | Reserved = 0 | | | PF=1 | Reserved = 0 | | | SP |
| 2<br>3<br>4<br>5<br>6 | Reserved = 0 | | | | | | | |
| 7<br>8 | (MSB) Parameter List Length (LSB) | | | | | | | |
| 9 | VU = 0 | | Reserved = 0 | | | | FLAG | LINK |

Figure 101. MODE SELECT (55)

The MODE SELECT (55) command provides a means for the initiator to specify LUN or device parameters to the Target. See the MODE SELECT (15) command for a description of the fields in this command.

# 7.11 PERSISTENT RESERVE IN (5E)

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | **7** | **6** | **5** | **4** | **3** | **2** | **1** | **0** |
| 0 | Command Code = 5Eh | | | | | | | |
| 1 | Reserved = 0 | | | Service Action | | | | |
| 2<br>3<br>4<br>5<br>6 | Reserved = 0 | | | | | | | |
| 7<br>8 | (MSB) Allocation Length (LSB) | | | | | | | |
| 9 | VU = 0 | | Reserved = 0 | | | | FLAG | LINK |

Figure 102. Persistent Reserve In (5E)

When a drive receives a PERSISTENT RESERVE IN command and RESERVE(6) or RESERVE(10) logical unit is active, the command is rejected with a RESERVATION CONFLICT status.

PERSISTENT RESERVE IN command does not conflict with a reservation established by the PERSISTENT RESERVE OUT command.

## 7.11.1 Service Action

The following service action codes are implemented. If a reserved service action code is specified, the drive returns a CHECK CONDITION status. The sense key is set to ILLEGAL REQUEST and the additional sense data is set to INVALID FIELD IN CDB.

| Code | Name | Descriptions |
|---|---|---|
| 00h | Read Keys | Reads all registered Reservation Keys |
| 01h | Read Reservations | Reads all current persistent reservations |
| 02h-1Fh | Reserved | Reserved |

Figure 103. Persistent Reserve In (5E) Service Action Codes

## 7.11.2 Parameter data for Read Keys

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | **7** | **6** | **5** | **4** | **3** | **2** | **1** | **0** |
| 0<br>\|<br>3 | (MSB)<br>Generation<br>(LSB) | | | | | | | |
| 4<br>\|<br>7 | (MSB)<br>Additional length (n-7)<br>(LSB) | | | | | | | |
| 8<br>\|<br>15 | (MSB)<br>First reservation key<br>(LSB) | | | | | | | |
| | : | | | | | | | |
| n-7<br>\|<br>n | (MSB)<br>Last reservation key<br>(LSB) | | | | | | | |

Figure 104. PERSISTENT RESERVE IN parameter data for Read Keys

**Generation** is a counter which increments when PERSISTENT RESERVE OUT command with “Register” or “Preempt and Clear” completes successfully. Generation is set to 0 as part of the power on reset process and hard reset process.

## 7.11.3 Parameter data for Read Reservations

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | **7** | **6** | **5** | **4** | **3** | **2** | **1** | **0** |
| 0<br>\|<br>3 | (MSB)<br>Generation<br>(LSB) | | | | | | | |
| 4<br>\|<br>7 | (MSB)<br>Additional length (n-7)<br>(LSB) | | | | | | | |
| 8<br>\|<br>n | (MSB)<br>Reservation descriptors<br>(LSB) | | | | | | | |

Figure 105. PERSISTENT RESERVE IN parameter data for Read Reservations

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | **7** | **6** | **5** | **4** | **3** | **2** | **1** | **0** |
| 0<br>\|<br>7 | (MSB)<br>Reservation key<br>(LSB) | | | | | | | |
| 8<br>\|<br>11 | (MSB)<br>Scope-specific address<br>(LSB) | | | | | | | |
| 12 | Reserved | | | | | | | |
| 13 | Scope | | | | Type | | | |
| 14<br>\|<br>15 | (MSB)<br>Extent length<br>(LSB) | | | | | | | |

Figure 106. PERSISTENT RESERVE IN Read Reservation Descriptor

**Scope-specific address** is filled with 0.

**Scope** and **Type** are described in PERSISTENT RESERVE OUT command section.

**Extent length** is filled with 0.

# 7.12 PERSISTENT RESERVE OUT (5F)

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | **7** | **6** | **5** | **4** | **3** | **2** | **1** | **0** |
| 0 | Command Code = 5Fh | | | | | | | |
| 1 | Reserved = 0 | | | | Service Action | | | |
| 2 | Scope | | | | Type | | | |
| 3<br>4<br>5<br>6 | Reserved = 0 | | | | | | | |
| 7<br>8 | (MSB) Parameter List Length = 18h (LSB) | | | | | | | |
| 9 | VU = 0 | | Reserved = 0 | | | | FLAG | LINK |

Figure 107. Persistent Reserve Out (5F)

When a drive receives a PERSISTENT RESERVE OUT command and RESERVE(6) or RESERVE(10) logical unit is active, the command is rejected with a RESERVE CONFLICT status.

**Parameter List Length** must be 18h. If not, the drive returns CHECK CONDITION status. And the sense key is set to ILLEGAL REQUEST and the additional sense data is set to PARAMETER LIST LENGTH ERROR

## 7.12.1 Service Action

The following service action codes are implemented. If a code which is not supported or a reserved code is specified, the drive returns a CHECK CONDITION status. The sense key is set to ILLEGAL REQUEST and the additional sense data is set to INVALID FIELD IN CDB. In case of PERSISTENT RESERVE OUT command executing a Register service action, this field is ignored.

| Code | Name | Description | Support |
|---|---|---|---|
| 00h | Register | Register a reservation key with the device server | Yes |
| 01h | Reserve | Create a persistent reservation using a reservation key | Yes |
| 02h | Release | Release a persistent reservation | Yes |
| 03h | Clear | Clear all reservation keys and all persistent reservations | No |
| 04h | Preempt | Preempt persistent reservations from another initiator | No |
| 05h | Preempt and Abort | Preempt persistent reservations from another initiator and clear the task set for the preempted initiator | Yes |
| 06h | Register and Ignor existing key | Register a reservation key with the drive | Yes |
| 07h - 1Fh | Reserved | Reserved | --- |

Figure 108. Persistent Reservation Service Action Code

**Register** This service action may conflict with a successfully established persistent reservation.

If the key specified in the Reservation Key field is not registered yet, a key specified in the Service Action Reservation Key is registered as new key. If the key is already registered but the initiator which registered the key is different from the initiator requesting the command, the drive returns RESERVATION CONFLICT status. If the key is already registered and the key is for the initiator requesting the command, the key is replaced with new key specified in the Service Action Reservation Key field.

Up to 4 keys can be held at the same time. When 4 keys are already registered, PERSISTENT RESERVE OUT command with Register service action which does not replace an existing key is rejected with RESERVATION CONFLICT status.

When PERSISTENT RESERVE OUT command with Register service action completes successfully, the Generation counter is incremented.

When keys are registered, the drive returns RESERVATION CONFLICT status against the RESERVE command and the RELEASE command.

**Reserve** This service action does not conflict with a successfully established persistent reservation.

If the initiator has not previously performed a Register service action, the command with this service action is rejected with RESERVATION CONFLICT status.

If the key specified in the Reservation key field is already registered but the initiator which registered the key is different from the initiator requesting the command, the target returns RESERVATION CONFLICT status.

If persistent reservation that is being attempted conflicts with persistent reservation that is held, the target returns a RESERVATION CONFLICT status.

The established persistent reservation applies to all commands received after the successful completion of the command.

**Release** This service action may conflict with a successfully established persistent reservation.

If the initiator requesting the command has not previously performed a Register service action, the command is rejected with a status of RESERVATION CONFLICT status.

If the key specified in the Reservation key field is already registered but the initiator which registered the key is different from the initiator requesting the command, the drive returns a RESERVATION CONFLICT status.

The drive returns GOOD status when a key specified in Reservation key field is not found.

When the key is found but the scope is different from the registered scope or the type is different from the registered type, the command is rejected with CHECK CONDITION status. The sense key is set to ILLEGAL REQUEST and the additional sense data is set to INVALID RELEASE OF ACTIVE PERSISTENT RESERVATION.

**Preempt and Abort** This service action does not conflict with a successfully established persistent reservation.

If the initiator requesting the command has not previously performed a Register service action, the command is rejected with a status of RESERVATION CONFLICT.

If the key specified in the Reservation key field is already registered but the initiator which registered the key is different from the initiator requesting the command, the drive returns a RESERVATION CONFLICT status.

Even if the key specified in the Service Action Reservation key field is not registered, the drive makes new persistent reservation without preempting if it does not conflict with an existing persistent reservation.

When the key is registered but reservation for the key is nothing yet and the new persistent reservation does not conflict with an existing persistent reservation, the drive makes new persistent reservation and clears the key and all commands from the initiator which was registered the key.

When the key specified in the Service Action Reservation Key is cleared even if a reservation for the key is nothing yet, the drive makes UNIT ATTENTION condition for the initiator which was registered the key. The sense key is set to UNIT ATTENTION and the additional sense data is set to RESERVATION PREEMPTED.

When PERSISTENT RESERVE OUT command with this service action completes successfully, the Generation counter is incremented.

## 7.12.2 Scope

Scope codes are implemented as follows. If a code which is not supported or a reserved code is specified, the drive returns a CHECK CONDITION status. The sense key is set to ILLEGAL REQUEST and the additional sense data is set to INVALID FIELD IN CDB. In case of PERSISTENT RESERVE OUT command executing a Register service action, this field is ignored.

| Code | Name | Description | Support |
|---|---|---|---|
| 0h | LU | Persistent reservation applies to the full logical unit | Yes |
| 1h | Extent | Persistent reservation applies to the specific extent | No |
| 2h | Element | Persistent reservation applies to the specific element | No |
| 3h - Fh | Reserved | Reserved | --- |

Figure 109. Persistent reserve Scope Code

## 7.12.3 Type

Type codes are implemented as follows. If a code which is not supported or a reserved code is specified, the target returns a CHECK CONDITION status. The sense key is set to ILLEGAL REQUEST and the additional sense data is set to INVALID FIELD IN CDB.

| Code | Name | Support | RD | WR | NWR | Add Rsv |
|---|---|---|---|---|---|---|
| 0h | Read Shared | No | --- | --- | --- | --- |
| 1h | Write Exclusive | Yes | SH | EX | PH | A1 |
| 2h | Read Exclusive | No | --- | --- | --- | --- |
| 3h | Exclusive Access | Yes | EX | EX | PH | A1 |
| 4h | Shared Access | No | --- | --- | --- | --- |
| 5h | Write Exclusive Registrants Only | Yes | SH | EO | PH | A2 |
| 6h | Exclusive Access Registrants Only | Yes | EO | EO | PH | A2 |
| 7h - Fh | Reserved | --- | --- | --- | --- | --- |

Figure 110. Persistent reserve Type Code

**RD** READ command and READ (10) command

**WR** WRITE command and WRITE (10) command

**NWR** Commands except following,

- READ command and READ (10) command
- WRITE command and WRITE (10) command
- RESERVE and RELEASE command

  If any key is registered, the target returns a RESERVATION CONFLICT status.
- PERSISTENT RESERVE IN command and PERSISTENT RESERVE OUT command

**SH** SHared: The target executes the command from all initiators.

**EX** EXclusive: The target executes the command from initiator which holds the the persistent reservation and rejects the command from the other initiators with RESERVATION CONFLICT status.

**PH** ProHibited: The target rejects the command from all initiators.

**EO** Exclusive registrant Only: The target executes the command from initiator which has registered a key and rejects the command from the other initiators with RESERVATION CONFLICT status.

**A1** The target rejects new PERSISTENT RESERVE on any types and returns RESERVATION CONFLICT status

**A2** The target accepts new PERSISTENT RESERVE with Write Exclusive Registrants Only type and Exclusive Access Registrants Only type. The target rejects new PERSISTENT RESERVE with Write Exclusive type and Exclusive Access type and returns RESERVATION CONFLICT status

| Exsisting Type / New Type | 1 | 3 | 5 | 6 |
|---|---|---|---|---|
| Write Exclusive (1) | Yes | Yes | Yes | Yes |
| Exclusive Access (3) | Yes | Yes | Yes | Yes |
| Write Exclusive Registrants Only (5) | Yes | Yes | No | No |
| Exclusive Access Registrants Only (6) | Yes | Yes | No | No |

Figure 111. Conflict between new and existing PERSISTENT RESERVE

## 7.12.4 Parameter list

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | **7** | **6** | **5** | **4** | **3** | **2** | **1** | **0** |
| 0 \| 7 | (MSB) Reservation Key (LSB) | | | | | | | |
| 8 \| 15 | (MSB) Service Action Reservation Key (LSB) | | | | | | | |
| 16 \| 19 | (MSB) Scope-specific address (LSB) | | | | | | | |
| 20 | Reserved | | | | | | | APTPL |
| 21 | Reserved | | | | | | | |
| 22 23 | (MSB) Extent length (LSB) | | | | | | | |

Figure 112. PERSISTENT RESERVE OUT parameter list

## 7.12.5 Summary

| Service Action | Parameters | | | | | | Gener-ation counter |
|---|---|---|---|---|---|---|---|
| | Scope Type | Rsv Key | SvcAct RsvKey | S-spec addr | Extent length | APTPL | |
| (0) Register | ignore | verify | save | ignore | ignore | apply | + 1 |
| (1) Reserve | apply | verify | ignore | ignore | ignore | ignore | --- |
| (2) Release | apply | verify | ignore | ignore | ignore | ignore | --- |
| (5) Preempt and Abort | apply | verify | save | ignore | ignore | ignore | + 1 |

Figure 113. Service Action and parameters and generation counter

### 7.12.5.1 Scope, Type

The Scope and the Type are applied in the process for the Reserve, Release, and Preempted and Clear service action but they are ignored in the process for the Register service action because they are not used.

### 7.12.5.2 Reservation Key

The Reservation Key is verified in each service action process. If the initiator which registered a key is different from the initiator requesting PERSISTENT RESERVE OUT command, the drive returns a RESERVATION CONFLICT status.

### 7.12.5.3 Service Action Reservation Key

On Register service action, the drive saves the key specified in the Service Action Reservation Key field as a key of initiator requesting PERSISTENT RESERVE OUT command.

On Preempt and Clear service action, the reservation which has a key specified in the Service Action Reservation Key field is preempted.

On other service actions, this field is ignored.

### 7.12.5.4 Scope-specified address

Parameter in the Scope-specified address field is ignored by the drive.

### 7.12.5.5 Extent length

Parameter in the Extent length field is ignored by the drive.

### 7.12.5.6 APTPL

The APTPL (Active Persist Through Power Loss) is valid only for the Register service action. The drive ignores the APTPL in other service actions.

The following table shows the relationship between the last valid APTPL value and information held by the drive.

| Information held by the drive | The last valid APTPL value | |
|---|---|---|
| | 0 | 1 |
| Registration | all keys are set to 0 | retained |
| Persistent Reservation | all are removed | retained |
| Generation counter | set to 0 | set to 0 |

Figure 114. APTPL and information held by a drive

### 7.12.5.7 Generation counter

The drive increments the Generation counter when Register service action or Preempt and Clear service action complete successfully.

# 7.13 PRE-FETCH (34)

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
| 0 | Command Code = 34h | | | | | | | |
| 1 | Reserved = 0 | | | Reserved = 0 | | | Immed = 0 | RelAdr = 0 |
| 2<br>3<br>4<br>5 | (MSB)<br>Logical Block Address<br>(LSB) | | | | | | | |
| 6 | Reserved = 0 | | | | | | | |
| 7<br>8 | (MSB)<br>Transfer Length<br>(LSB) | | | | | | | |
| 9 | VU = 0 | | Reserved = 0 | | | | FLAG | LINK |

Figure 115. Pre-Fetch (34)

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

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

**Immed** (Immediate) must be zero. An immediate bit of zero indicates that the status shall not be returned until the operation has completed.

If the Immed bit is set to one, the drive returns a Check Condition status. The sense key shall be set to Illegal Request and the additional sense code shall be set to Invalid Field in CDB.

**RelAdr** Relative Block Address is not supported. Must be set to zero.

# 7.14 READ (6) - (08)

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | **7** | **6** | **5** | **4** | **3** | **2** | **1** | **0** |
| 0 | Command Code = 08h | | | | | | | |
| 1 | Reserved = 0 | | | (MSB) | | | LBA | |
| 2<br>3 | Logical Block Address<br>(LSB) | | | | | | | |
| 4 | Transfer Length | | | | | | | |
| 5 | VU = 0 | | Reserved = 0 | | | | FLAG | LINK |

Figure 116. READ (6) - (08)

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

**Logical block address** This field specifies the logical unit at which the read operation shall begin.

**Transfer length** This field specifies the number of blocks to be transferred. A value of zero implies 256 blocks are to be transferred.

**Note:** Errors are handled by ERP (error recovery procedure). ERP's are controlled by the error recovery parameters specified by MODE SELECT command.

# 7.15 READ (10) - (28)

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | **7** | **6** | **5** | **4** | **3** | **2** | **1** | **0** |
| 0 | Command Code = 28h | | | | | | | |
| 1 | Reserved = 0 | | | DPO | FUA | Reserved = 0 | | RelAdr = 0 |
| 2<br>3<br>4<br>5 | (MSB)<br>Logical Block Address<br>(LSB) | | | | | | | |
| 6 | Reserved = 0 | | | | | | | |
| 7<br>8 | (MSB) Transfer Length<br>(LSB) | | | | | | | |
| 9 | VU = 0 | | Reserved = 0 | | | | FLAG | LINK |

Figure 117. Read (10) - (28)

The READ (10) command requests the drive to transfer data to the initiator. The larger Logical Block Address and Transfer Length fields permit greater quantities of data to be requested per command than with the READ command and are required to access the full LBA range of the larger capacity drives.

**Transfer length** The number of contiguous blocks to be transferred. If the transfer length is zero, the seek occurs, but no data is transferred. This condition is not considered an error. If read ahead is enabled, a read ahead is started after the seek completes.

**DPO** A DPO (Disable Page Out) bit of one indicates that the data accessed by this command is to be assigned the lowest priority for being written into or retained by the cache. A DPO bit of one overrides any retention priority specified in the Mode Select Page 8 Caching Parameters. A DPO bit of zero indicates the priority is determined by the retention priority. The initiator should set the DPO bit when the blocks read by this command are not likely to be read again in the near future.

**FUA** Force Unit Access. A FUA bit of one indicates that the data is read from the media and not from the cache. A FUA bit of zero allows the data to be read from either the media or the cache.

**RelAdr** Relative Block Address is not supported. Must be set to zero.

# 7.16 READ BUFFER (3C)

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
| 0 | Command Code = 3Ch | | | | | | | |
| 1 | Reserved = 0 | | | | Mode | | | |
| 2 | Buffer ID = 0 | | | | | | | |
| 3<br>4<br>5 | (MSB)<br>Buffer Offset<br>(LSB) | | | | | | | |
| 6<br>7<br>8 | (MSB)<br>Allocation Length<br>(LSB) | | | | | | | |
| 9 | VU = 0 | | Reserved = 0 | | | | FLAG | LINK |

Figure 118. READ BUFFER (3C)

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

The function of this command and the meaning of fields within the command descriptor block depend on the contents of the mode field.

| MODE | Description |
|---|---|
| **0000** | Read Combined Header and Data |
| **0010** | Data |
| **0011** | Descriptor |
| **1010** | Read Data from Echo Buffer |
| **1011** | Echo Buffer Descriptor |
| **All others** | Not supported. |

## 7.16.1 Combined Header And Data (Mode 0000b)

In this mode a four byte header followed by data bytes are returned to the initiator during the DATA IN phase. The Buffer ID and the buffer offset field are reserved.

The drive terminates the DATA IN phase when allocation length bytes of header plus data have been transferred or when the header and all available data have been transferred to the initiator, whichever is less.

The four-byte READ BUFFER header (see figure below) is followed by data bytes from the data buffer of the drive.

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
| 0 | RSVD = 0 | | | | | | | |
| 1<br>2<br>3 | (MSB) Buffer Capacity<br><br>(LSB) | | | | | | | |

Figure 119. READ BUFFER Header

The buffer capacity specifies the total number of data bytes that are available in the data buffer of the drive. This number is not reduced to reflect the allocation length nor is it reduced to reflect the actual number of bytes written using the WRITE BUFFER command.

Following the READ BUFFER header the drive will transfer data from its data buffer.

## 7.16.2 Read Data (Mode 0010b)

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

**Buffer ID** field must be set to zero, indicating the data transfer buffer. If another value is specified, the command is terminated with CHECK CONDITION status. The drive shall set sense key to ILLEGAL REQUEST and additional sense code to ILLEGAL FIELD IN CDB.

**Buffer Offset** specifies the offset of the memory space specified by the Buffer ID. The initiator should conform to the offset boundary requirements returned in the READ BUFFER descriptor. If the value exceeds the buffer specified, the command is terminated with CHECK CONDITION status. The drive shall set sense key to ILLEGAL REQUEST and additional sense code to ILLEGAL FIELD IN CDB.

**Allocation Length** The drive terminates the DATA IN phase when allocation length bytes of data have been transferred or when the header and all available data have been transferred to the initiator, whichever is less.

## 7.16.3 Descriptor (Mode 0011b)

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

**Buffer ID** field should normally be set to zero indicating the drive data transfer buffer. If any other value is specified, the drive returns all zeros in the READ BUFFER descriptor.

**Buffer Offset** field is reserved.

**Allocation Length** should be set to four or greater. The drive transfers the allocation length or four bytes of READ BUFFER descriptor, whichever is less. The READ BUFFER descriptor is defined in the figure below.

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | **7** | **6** | **5** | **4** | **3** | **2** | **1** | **0** |
| 0 | Offset Boundary | | | | | | | |
| 1<br>2<br>3 | (MSB) Buffer Capacity<br><br>(LSB) | | | | | | | |

Figure 120. Read Buffer Descriptor

The value contained in the Buffer Offset field of subsequent WRITE BUFFER and READ BUFFER commands should be a multiple of two to the power of the offset boundary. The offset boundary is always set to nine, which indicates Sector Boundaries.

## 7.16.4 Read Data from Echo Buffer (Mode 1010b)

In this mode the drive transfers data from the echo buffer. The echo buffer will transfer the same data as when the Write Buffer command was issued with the mode field set to echo buffer.

Write Buffer command with the mode field set to echo buffer should be sent prior to the Read Buffer command; otherwise the Read Buffer command will be terminated with Check Condition "Illegal Request".

## 7.16.5 Echo Buffer Descriptor (Mode 1011b)

In this mode, a maximum of four bytes of Read Buffer Descriptor information is returned. The drive returns the descriptor information for the echo buffer. The Buffer Offset field is reserved in this mode and must be zero. The drive transfers the lesser of the allocation length or four bytes of following Echo Buffer Descriptor.

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | **7** | **6** | **5** | **4** | **3** | **2** | **1** | **0** |
| 0 | Reserved = 0 | | | | | | | EBOS=<br>0 |
| 1 | Reserved = 0 | | | | | | | |
| 2 | Reserved = 0 | | | | (MSB) Buffer Capacity | | | |
| 3 | Buffer Capacity (LSB) | | | | | | | |

Figure 121. Echo Buffer Descriptor

**EBOS** (Echo Buffer Overwritten Supported) bit of zero indicates that the echo buffer is shared by all initiators.

**Buffer Capacity** field returns the size of the echo buffer in byte aligned to a four-byte boundary.

# 7.17 READ CAPACITY (25)

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
| 0 | Command Code = 25h | | | | | | | |
| 1 | Reserved = 0 | | | Reserved = 0 | | | | RelAdr = 0 |
| 2<br>3<br>4<br>5 | (MSB)<br>Logical Block Address<br>(LSB) | | | | | | | |
| 6<br>7 | Reserved = 0 | | | | | | | |
| 8 | Reserved = 0 | | | | | | | PMI |
| 9 | VU = 0 | | Reserved = 0 | | | | FLAG | LINK |

Figure 122. READ CAPACITY (25)

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

- **Logical Block Address** is used in conjunction with the PMI bit.
- **RelAdr** A Relative Address is not supported. Must be set to zero.
- **PMI** Partial Medium Indicator indicates

| PMI | Description |
|---|---|
| 0 | The drive returns the last logical block address of the drive. |
| 1 | The drive returns the last logical block address and block length in bytes are that of the logical block address after which a substantial delay in data transfer will be encountered. This returned logical block address shall be greater than or equal to the logical block address specified by the logical block address fields in the command descriptor block.<br><br>This option provides the information which the initiator needs to determine the amount of space available on the same track which is accessible without a head switch or seek. |

### 7.17.1 Returned Data Format

The data returned to the initiator in response to the READ CAPACITY command is described here. The data is returned in the DATA IN phase.

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | **6** | **7** | **5** | **4** | **3** | **2** | **1** | **0** |
| 0<br>1<br>2<br>3 | (MSB)<br>Logical Block Address<br>(LSB) | | | | | | | |
| 4<br>5<br>6<br>7 | (MSB)<br>Block Length<br>(LSB) | | | | | | | |

Figure 123. Format of READ CAPACITY command reply

- **Block Length** specifies the length in bytes of the block.

# 7.18 READ DEFECT DATA (37)

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
| 0 | Command Code = 37h | | | | | | | |
| 1 | Reserved = 0 | | | Reserved = 0 | | | | 0 |
| 2 | Reserved = 0 | | | Plist | Glist | Defect List Format | | |
| 3<br>4<br>5<br>6 | Reserved = 0 | | | | | | | |
| 7<br>8 | (MSB) Allocation Length (LSB) | | | | | | | |
| 9 | VU = 0 | | Reserved = 0 | | | | FLAG | LINK |

Figure 124. Read Defect Data (37)

The READ DEFECT DATA command requests that the Target transfer the medium defect data to the initiator.

If the target is unable to access any medium defect data it will return a Check Condition status with the appropriate sense key. The sense key will be set to either Medium Error(03h) if a medium error occurred or No Sense(00h) if the list does not exist and the additional sense code will be set to Defect List Error (19h).

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

**Glist** The Grown Defect List (Glist) bit set to one indicates that the target returns the grown defect list. A Glist bit of zero indicates that the target shall not return the Grown Defect list of defects.

**Note:** With both bits set to one Plist and Glist the target will return both the Primary and Grown defect lists. With both bits set to zero, the target will return only a four-byte Defect List Header.

**Defect List format**

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

The Defect List Format of '100 (Bytes from Index Format)' and '101 (Physical Sector Format)' are supported. If the requested format is not supported by the drive, it will return the defect list in its default format '101' and then terminate the command with Check Condition status. The sense key will be set to Recovered Error (01h) and the additional sense code will be set to Defect List Not Found(1Ch).

The drive sends defect list (Defect Descriptors) in a four byte ABA (Absolute Block Address) format which follows a four byte Defect List Header.

The target will transfer all of the Read Defect Data up to the number of bytes allocated by the initiator.

**Note:** The drive will terminate the Data In phase when the Allocation Length has been transferred or when all available Defect Data has been transferred to the initiator, whichever is less.

The Read Defect Data contains a four byte header followed by zero or more defect descriptors.

## 7.18.1 Defect List Header

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
| | Defect List Header | | | | | | | |
| 0 | Rsvd = 0 | | | | | | | |
| 1 | Reserved = 0 | | | | Plist | Glist | Defect List Format | |
| 2<br>3 | (MSB) Defect List length (LSB) | | | | | | | |

Figure 125. Defect List Header

## 7.18.2 Bytes from Index Format (100b)

| Byte | Defect Descriptors |
|---|---|
| 0<br>1<br>2 | (MSB)<br>Cylinder Number of Defect<br>(LSB) |
| 3 | Head Number of Defect |
| 4<br>5<br>6<br>7 | (MSB)<br>Defect Bytes from Index<br>(LSB) |

Figure 126. Defect Descriptors of Bytes from Index Format

Defect Bytes from Index is gotten using the following equation:

```
Bytes from Index = (Physical Sector Number) * N
Where: N = Bytes per sector
```

## 7.18.3 Physical Sector Format (101b)

| Byte | Defect Descriptors |
|---|---|
| 0<br>1<br>2 | (MSB)<br>Cylinder Number of Defect<br>(LSB) |
| 3 | Head Number of Defect |
| 4<br>5<br>6<br>7 | (MSB)<br>Defective Sector Number<br>(LSB) |

Figure 127. Defect Descriptors of Physical Sector Format

The Defect List Format field specifies the format of the defect list data returned by the target.

The Defect List Length field specifies the length in bytes of the defect descriptors that follow. The Defect List Length is equal to eight times the number of defect descriptors.

Normally the Target will set the Defect List Length field to the amount of space needed to contain the entire defect list. However, the Target is capable of building a defect list with a length such that the entire list cannot be transferred using the maximum allocation length. If the defect list grows beyond 8191 entries, the defect data cannot be transferred with an allocation length of 0FFFFh. The Target will transfer a partial defect list and return Check Condition status with the sense key set to Recovered Error and the Additional Sense Code is set to Partial Defect List Transferred. The defect list length will be set to 0FFF8h, indicating the maximum number of defect descriptors which can be transferred. Defects beyond this number can not be read by the initiator.

# 7.19 READ DEFECT DATA (B7)

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
| 0 | Command Code = B7h | | | | | | | |
| 1 | Reserved = 0 | | | Plist | Glist | Defect List Format | | |
| 2<br>3<br>4<br>5 | Reserved = 0 | | | | | | | |
| 6<br>\|<br>9 | (MSB)<br>Allocation Length<br>(LSB) | | | | | | | |
| 10 | Reserved = 0 | | | | | | | |
| 11 | VU = 0 | | Reserved = 0 | | | | FLAG | LINK |

Figure 128. Read Defect Data (B7)

The READ DEFECT DATA command requests that the Target transfer the medium defect data to the initiator.

If the target is unable to access any medium defect data it will return a Check Condition status with the appropriate sense key. The sense key will be set to either Medium Error (03h) if a medium error occurred or to No Sense (00h) if the list does not exist and the additional sense code will be set to Defect List Error (19h).

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

**Glist** The Grown Defect List (Glist) bit set to one indicates that the target returns the grown defect list. A Glist bit of zero indicates that the target shall not return the Grown Defect list of defects.

**Note:** With both bits set to one Plist and Glist the target will return both the Primary and Grown defect lists. With both bits set to zero the target will return only a four-byte Defect List Header.

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

The Defect List Format of '100 (Bytes from Index Format)' and '101 (Physical Sector Format)' are supported. If the requested format is not supported by the drive, it will return the defect list in its default format '101' and then terminate the command with Check Condition status. The sense key will be set to Recovered Error (01h) and the additional sense code will be set to Defect List Not Found (1Ch).

The drive sends defect list (Defect Descriptors) in a four byte ABA (Absolute Block Address) format which follows a four byte Defect List Header.

The target will transfer all of the Read Defect Data up to the number of bytes allocated by the initiator.

**Note:** The drive will terminate the Data In phase when the Allocation Length has been transferred or when all available Defect Data has been transferred to the initiator, whichever is less.

The Read Defect Data contains a four byte header, followed by zero or more defect descriptors.

## 7.19.1 Defect List Header

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
| | **Defect List Header** | | | | | | | |
| 0 | Rsvd = 0 | | | | | | | |
| 1 | Reserved = 0 | | | | Plist | Glist | Defect List Format | |
| 2<br>3 | Rsvd = 0 | | | | | | | |
| 4<br>\|<br>7 | (MSB)<br>Defect List length<br>(LSB) | | | | | | | |

Figure 129. Defect List Header

## 7.19.2 Bytes from Index Format (100b)

| Byte | Defect Descriptors |
|---|---|
| 0<br>1<br>2 | (MSB)<br>Cylinder Number of Defect<br>(LSB) |
| 3 | Head Number of Defect |
| 4<br>5<br>6<br>7 | (MSB)<br>Defect Bytes from Index<br>(LSB) |

Figure 130. Defect Descriptors of Bytes from Index Format

Defect Bytes from Index is derived using the following equation:

Bytes from Index = (Physical Sector Number) + N

where N = Bytes per sector

## 7.19.3 Physical Sector Format (101b)

| Byte | Defect Descriptors |
|---|---|
| 0<br>1<br>2 | (MSB)<br>Cylinder Number of Defect<br>(LSB) |
| 3 | Head Number of Defect |
| 4<br>5<br>6<br>7 | (MSB)<br>Defective Sector Number<br>(LSB) |

Figure 131. Defect Descriptors of Physical Sector Format

The Defect List Format field specifies the format of the defect list data returned by the Target.

The Defect List Length field specifies the length in bytes of the defect descriptors that follow. The Defect List Length is equal to eight times the number of defect descriptors.

If the Allocation Length is insufficient to transfer all of the defect descriptors, the Defect List Length will not be adjusted to reflect the truncation. The Target will not create a CHECK CONDITION status.

# 7.20 READ LONG (3E)

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
| 0 | Command Code = 3Eh | | | | | | | |
| 1 | Reserved = 0 | | | Reserved = 0 | | | CORT = 0 | RelAdr = 0 |
| 2<br>3<br>4<br>5 | (MSB)<br>Logical Block Address<br>(LSB) | | | | | | | |
| 6 | Reserved = 0 | | | | | | | |
| 7<br>8 | (MSB)<br>Byte Transfer Length<br>(LSB) | | | | | | | |
| 9 | VU = 0 | | Reserved = 0 | | | | FLAG | LINK |

Figure 132. READ LONG (3E)

The READ LONG command requests the drive to transfer one block of data to the initiator. The transfer data includes data and ECC field data.

- **CORT** bit of ZERO causes the logical block to be read without any correction made by the drive. A CORT bit of one is not supported by the Target. (A corrected bit of one causes the data to be corrected by ECC before transferring the data to the initiator.)
- **Logical Block Address** field specifies the logical block at which the read operation shall occur.
- **Byte Transfer Length**. This field must specify exactly the number of bytes of data that are available for transfer. If a non-zero byte transfer length does not match the available data length, the target terminates the command with CHECK CONDITION status, the sense key is set to ILLEGAL REQUEST, and an additional sense code set to INVALID FIELD IN CDB. The valid and ILI bits is set to one and the information field is set to the difference of the requested length minus the actual length in bytes. Negative values are indicated by two's complement notation.

  The transfer length is calculated as follows:

  *transfer length = logical block size + 40*
- **RelAdr** Relative Block Address is not supported by the drive.

# 7.21 REASSIGN BLOCKS (07)

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | **7** | **6** | **5** | **4** | **3** | **2** | **1** | **0** |
| 0 | Command Code = 07h | | | | | | | |
| 1 | Reserved = 0 | | | | Reserved = 0 | | | |
| 2<br>3<br>4 | Reserved = 0 | | | | | | | |
| 5 | VU = 0 | | Reserved = 0 | | | | FLAG | LINK |

Figure 133. REASSIGN BLOCKS (07)

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

Reassignment is complete upon the completion of the REASSIGN BLOCKS command. At this time, the defective logical block address has been added to the grown (“G” list) defect list.

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

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

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | **7** | **6** | **5** | **4** | **3** | **2** | **1** | **0** |
| 0 | RSVD = 0 | | | | | | | |
| 1 | RSVD = 0 | | | | | | | |
| 2<br>3 | (MSB) Defect list length = 4/8/12/16 (LSB) | | | | | | | |
| 4<br>5<br>6<br>7 | (MSB) Defect Logical Block Address –1 (LSB) | | | | | | | |
| 8<br>9<br>10<br>11 | (MSB) Defect Logical Block Address –2 (LSB) | | | | | | | |
| 12<br>13<br>14<br>15 | (MSB) Defect Logical Block Address –3 (LSB) | | | | | | | |
| 16<br>17<br>18<br>19 | (MSB) Defect Logical Block Address –4 (LSB) | | | | | | | |

Figure 134. Format of REASSIGN BLOCKS data

- **Defect List Length** must be 4,8,12, or 16. Otherwise, the drive returns Check Condition with Sense key = Illegal request.
- **Defective Logical Block Address** is 4 bytes in length. The initiator can specify from 1 to 4 Defective Logical Block Address according to the Defect List Length from 4 to 16, respectively. Defective Logical Block Addresses must be ordered in ascending order, otherwise the drive returns Check Condition.

# 7.22 RECEIVE DIAGNOSTICS (1C)

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
| 0 | Command Code = 1Ch | | | | | | | |
| 1 | Reserved = 0 | | | Reserved = 0 | | | | PCV |
| 2 | Page Code | | | | | | | |
| 3<br>4 | (MSB) Parameter List Length (LSB) | | | | | | | |
| 5 | VU = 0 | | Reserved = 0 | | | | FLAG | LINK |

Figure 135. RECEIVE DIAGNOSTIC RESULTS (1C)

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

- **PCV** (Page Code Valid) bit of zero indicates that the most recent SEND DIAGNOSTIC command shall define the data returned by this command. PCV bit of one indicates that the contents of the Page Code field shall define the data returned by this command.
- **Parameter List Length** specifies the amount of data to be returned to the initiator. This value may be zero and this is not considered an error. The target terminates the Data In phase when all available data has been transferred or when the number of bytes transferred equals the Parameter List Length.

## 7.22.1 Receive Diagnostic Results Page 0

This page contains a list of supported pages.

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
| 0 | Page Code = 0 | | | | | | | |
| 1 | Reserved = 0 | | | | | | | |
| 2-3 | Page Length = 03h | | | | | | | |
| 4 | First supported page 0h | | | | | | | |
| 5 | Second supported page 40h | | | | | | | |
| 6 | Third supported page 80h | | | | | | | |

Figure 136. Receive Diagnostic Results page 0

## 7.22.2 Receive Diagnostic Results Page 40

Using the Send diagnostic command, an address in either physical or logical format is supplied to the drive. This page is then used to retrieve the address translated into the other format.

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
| 0 | Page Code = 40h | | | | | | | |
| 1 | Reserved = 0 | | | | | | | |
| 2-3 | Page Length = 0Ah | | | | | | | |
| 4 | Reserved = 0 | | | | | Supplied format | | |
| 5 | RA | ALTS | ALTT | Reserved=0 | | Translate format | | |
| 6 - 13 | Translated Address | | | | | | | |

Figure 137. Receive Diagnostic Results Page 40

- **Supplied Format** is the value supplied by the Send Diagnostic command; it may be one of the three following values:
  - **000b** Block format
  - **100b** Bytes From Index format
  - **101b** Physical Sector format

  It specifies the format in which the address has been supplied.
- **Translate Format** is the value supplied by the Send Diagnostic command and specifies the format in which the address has been translated into List. If the supplied format is the Block format, the Translate format must be either Bytes from Index or Physical Sector format. If the supplied format is the Bytes from Index or Physical Sector format, the Translate format must be Block format. Otherwise the Target will terminate the command with Check Condition status.
- **RA (Reserved Area)** is set to one if the translated block is a reserved area.
- **ALTS (Alternate Sector)** is set to one if the translated block is in alternate sector area.
- **ALTT (Alternate Track)** is set to one if the translated block is in alternate track area. This bit is not used by the drive.

- **Translated Address** contains the address in the translate format. If it is a logical block address, it is contained within the first 4 bytes of the field (bytes 6 to 9) of the page data. For a physical format it is as follows:

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
| 6-8 | Cylinder Number | | | | | | | |
| 9 | Head Number | | | | | | | |
| 0-13 | Sector Number or Bytes from Index | | | | | | | |

Figure 138. Translated address

## 7.22.3 Receive Diagnostic Results Page 80

This page contains the off-line read scan status.

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
| 0 | Page Code = 80h | | | | | | | |
| 1 | Reserved = 0 | | | | | | | |
| 2-3 | Page Length = 04h | | | | | | | |
| 4 | SMART Revision = 03h | | | | | | | |
| 5 | Off-line Read Scan Ratio | | | | | | | |
| 6<br>7 | (MSB) Estimated Completion Time | | | | | | | |

Figure 139. Receive Diagnostic Results Page 80

- **Off-line Read Scan Ratio** indicates the progress of the read scan activity. The number is the numerator of a fraction whose denominator is 100.
- **Estimated Completion Time** indicates the total time in seconds to complete a full read scan without interruption.

# 7.23 RELEASE (17)

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
| 0 | Command Code = 17h | | | | | | | |
| 1 | Reserved = 0 | | | 3rdPty | 3rd Party ID | | | Ext=0 |
| 2 | Reservation Identification | | | | | | | |
| 3<br>4 | Reserved = 0 | | | | | | | |
| 5 | VU = 0 | | Reserved = 0 | | | | FLAG | LINK |

Figure 140. RELEASE (17)

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

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

- **3rdPty** bit indicates that
  - **1** this release process is for a third party which is specified by 3rd Party ID.
  - **0** this release process is for the initiator itself.
- **3rd Party ID** specifies the ID of the third party for which the LUN is reserved. (Refer to Section 7.28**, RESERVE (16)** on page 152)
- **Ext**ents must be 0. Extension is not supported by the drive.
- **Reservation Identification** field is ignored.

# 7.24 RELEASE (57)

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | **7** | **6** | **5** | **4** | **3** | **2** | **1** | **0** |
| 0 | Command Code = 57h | | | | | | | |
| 1 | Reserved = 0 | | | 3rdPty | Reserved = 0 | | | Ext = 0 |
| 2 | Reservation Identification | | | | | | | |
| 3 | 3rd Party Device ID | | | | | | | |
| 4 | Reserved = 0 | | | | | | | |
| 5 | Reserved = 0 | | | | | | | |
| 6 | Reserved = 0 | | | | | | | |
| 7 | Reserved = 0 | | | | | | | |
| 8 | Reserved = 0 | | | | | | | |
| 9 | VU = 0 | | Reserved = 0 | | | | FLAG | LINK |

Figure 141. RELEASE (57)

The Release command is used to release a LUN previously reserved. 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 the reservation.

Extents are not supported by the drive. The Ext (Extent) bit must be zero and the Reservation Identification field is ignored. If the Ext bit is not zero, Check Condition Status is returned with a Sense Key of Illegal Request and additional sense code of Invalid Field in CDB.

If the 3rdPty bit is zero, then the third-party release option is not requested. If the 3rdPty bit is one, then the Target releases the LUN, but only if the reservation was made using the third-party reservation option and the 3rd Party ID is the ID of the initiator that made the reservation

# 7.25 REPORT DEVICE IDENTIFIER (A3)

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | **7** | **6** | **5** | **4** | **3** | **2** | **1** | **0** |
| 0 | Command Code = A3h | | | | | | | |
| 1 | Reserved = 0 | | | | Service Action = 05h | | | |
| 2 | Reserved = 0 | | | | | | | |
| 3 | Reserved = 0 | | | | | | | |
| 4<br>5 | (MSB) LUN (LSB) | | | | | | | |
| 6<br>\|<br>9 | (MSB) Allocation Length (LSB) | | | | | | | |
| 10 | Reserved = 0 | | | | | | | |
| 11 | VU = 0 | | Reserved = 0 | | | | FLAG | LINK |

Figure 142. REPORT DEVICE IDENTIFIER(A3)

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

The LUN contains the logical unit number. If the requested logical unit has not been added to the target, the command is terminated with a CHECK CONDITION status. The sense key is set to ILLEGAL REQUEST, and the additional sense code set to LOGICAL UNIT NOT SUPPORTED.

The ALLOCATION LENGTH field indicates how much space has been reserved for the returned parameter data. If the length is not sufficient to contain all the parameter data, the first portion of the data is returned. This is not considered an error. The actual length of the parameter data is available in the IDENTIFIER LENGTH field in the parameter data. If the remainder of the parameter data is required, the application client should send a new REPORT DEVICE IDENTIFIER command with an ALLOCATION LENGTH field large enough to contain all the data.

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

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | **7** | **6** | **5** | **4** | **3** | **2** | **1** | **0** |
| 0<br>\|<br>3 | (MSB) Identifier Length = n-3 (LSB) | | | | | | | |
| 4<br>n | Identifier | | | | | | | |

Figure 143. REPORT DEVICE IDENTIFIER parameter list

The IDENTIFIER LENGTH field specifies the length in bytes of the IDENTIFIER field. If the ALLOCATION LENGTH field in the CDB is too small to transfer all of the identifier, the length is not adjusted to reflect the truncation. The identifier length initially equals zero, and is changed only by a successful SETDEVICE IDENTIFIER command.

The IDENTIFIER field contains a vendor specific value. The value reported is the last value written by a successful SET DEVICE IDENTIFIER command. The value of the identifier is changed only by a SET DEVICE IDENTIFIER command. The identifier value persist through resets, power cycles, media format operations.

The target return the same Identifier to all initiators on all ports.

The execution of a REPORT DEVICE IDENTIFIER requires the enabling of a nonvolatile memory within the logical unit. If the nonvolatile memory is not ready, the device server return CHECK CONDITION status, rather than wait for the device to become ready. The sense key is set to NOT READY and the additional sense data is set as described in the TEST UNIT READY command. This information should allow the application client to determine the action required to cause the device server to become ready.

# 7.26 REPORT LUNS (A0)

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | **7** | **6** | **5** | **4** | **3** | **2** | **1** | **0** |
| 0 | Command Code = A0h | | | | | | | |
| 1<br>\|<br>5 | Reserved | | | | | | | |
| 6<br>\|<br>9 | (MSB)<br>Allocation Length<br>(LSB) | | | | | | | |
| 10 | Reserved | | | | | | | |
| 11 | VU = 0 | | Reserved = 0 | | | | FLAG | LINK |

Figure 144. REPORT LUNS (A0)

The Report LUN's command requests that the target return the known Logical Unit Numbers (LUN) to the initiator. The Report LUN's command should always be available and is unaffected by any reservations.

The Allocation Length must be at least 16 bytes. If the Allocation Length is less than 16 bytes, the target will return in Check Condition Status with Sense Key of Illegal Request and additional Sense Code of Invalid Field in CDB. If the Allocation Length is not sufficient to contain the Logical Unit Number values for all configured logical units, the target shall report as many logical unit number values as will fit in the specified Allocation Length. This is not considered an error.

The Report LUN's command will send the LUN List in the subsequent Data Out Phase. The format of the LUN List is shown in the following table.

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | **7** | **6** | **5** | **4** | **3** | **2** | **1** | **0** |
| 0<br>\|<br>3 | (MSB)<br>LUN List Length = 8<br>(LSB) | | | | | | | |
| 4<br>\|<br>7 | Reserved | | | | | | | |
| 8<br>\|<br>15 | (MSB)<br>LUN = 0<br>(LSB) | | | | | | | |

Figure 145. LUN Reporting Parameter List Format

The LUN list length shall contain the length in bytes of the LUN list that is available to be transferred. This product only supports one LUN. Therefore, the LUN list length must be set to 8. The only supported LUN is zero.

# 7.27 REQUEST SENSE (03)

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
| 0 | Command Code = 03h | | | | | | | |
| 1 | Reserved = 0 | | | Reserved = 0 | | | | |
| 2<br>3 | Reserved = 0 | | | | | | | |
| 4 | Allocation Length | | | | | | | |
| 5 | VU = 0 | | Reserved = 0 | | | | FLAG | LINK |

Figure 146. REQUEST SENSE (03)

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

The sense data shall be available under the following conditions:

- The previous command to the specified I_T_L[4] nexus terminated with CHECK CONDITION status. An other information (e.g. medium position) is available in any fields.
- The previous command to the specified I_T_L nexus ended unexpected BUS FREE error.

If REQUEST SENSE command with an invalid LUN is received, the drive returns GOOD status and reports a sense key of ILLEGAL REQUEST and an additional sense code of LOGICAL UNIT NOT SUPPORTED.

If the drive has no sense data available to return, it shall return a sense key of NO SENSE and an additional sense code of NO ADDITIONAL SENSE INFORMATION.

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

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

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

The contents of the sense data are defined in Section 11.0, "SCSI Sense Data" on page 231

[4] A nexus which exists between an initiator, a target, and a logical unit.

# 7.28 RESERVE (16)

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
| 0 | Command Code = 16h | | | | | | | |
| 1 | Reserved = 0 | | | 3rdPty | 3rd Party ID | | | Ext = 0 |
| 2 | Reservation Identification | | | | | | | |
| 3<br>4 | (MSB) Extent List Length = 0 (LSB) | | | | | | | |
| 5 | VU = 0 | | Reserved = 0 | | | | FLAG | LINK |

Figure 147. RESERVE (16)

The RESERVE command is used to reserve a LUN for an initiator. This reservation can be either for the initiator which sends this command or for the third party which is specified in this command.

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

- The reservation is superseded by another valid RESERVE command from the initiator that made the reservation.
- The LUN is released by a RELEASE command from the same initiator.
- A hard reset condition occurs. (A SCSI bus Reset assertion).
- A TARGET RESET message is received from any initiator.
- Power off/on occurs.

**3rdPty** bit is to indicates that

**1** this reservation is for a third party which is specified by 3rd Party ID.
**0** this reservation is for the initiator itself.

**3rd Party ID** specifies the ID of the third party for which the LUN is reserved.

**Note:** The LUN may be only released by the initiator who sent the RESERVE command.

**Ext**ents must be 0. Extension is not supported by the drive.

**Reservation Identification** is ignored.

**Extent List** length must be zero. Extent List length is not supported.

# 7.29 RESERVE (56)

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | **7** | **6** | **5** | **4** | **3** | **2** | **1** | **0** |
| 0 | Command Code = 57h | | | | | | | |
| 1 | Reserved = 0 | | | 3rdPty | Reserved | | | Ext = 0 |
| 2 | Reservation Identification | | | | | | | |
| 3 | Third Pary Device ID | | | | | | | |
| 4 | Reserved = 0 | | | | | | | |
| 5 | Reserved = 0 | | | | | | | |
| 6 | Reserved = 0 | | | | | | | |
| 7<br>8 | (MSB) Extent List Length = 0 (LSB) | | | | | | | |
| 9 | VU = 0 | | Reserved = 0 | | | | FLAG | LINK |

Figure 148. RESERVE (56)

The Reserve command is used to reserve a LUN for an Initiator. This reservation can be either for the Initiator sending the command or for a third party as specified by the Initiator.

Extents are not supported by the drive. The Ext bit must be zero. If Ext bit is set to one, Check Condition Status is returned with a Sense Key of Illegal Request and additional sense code of Invalid Field in CDB The Reservation Identification and Extent List Length fields are ignored.

The Reserve command requests that the entire Lun be reserved for the Initiator until

- the reservation is superseded by another valid Reserve command from the Initiator that made the reservation.
- the reservation is released by a Release command from the same Initiator.
- a hard Reset condition occurs.
- a Target Reset message is received from any Initiator.
- a power off/on cycle occurs.

The 3rdPty bit of zero indicates that the Initiator that issued the Reserve command is the Initiator for which the LUN is reserved.

The 3rdPty bit of one indicates that this is a third-party reservation. The 3rd Party ID byte specifies the ID of the third party for which the LUN is reserved. A reservation made with the 3rdPty bit of one and the 3rd Party ID byte set to the Initiator that issued this Reserve command is considered equivalent to a reservation made with the 3rdPty bit set to zero.

Only the Initiator that issued the Reserve command for a LUN may release the LUN, regardless of the 3rdPty option. This Initiator may also release the LUN by issuing another Reserve command. This superseding Reserve command releases the previous reservation when the new reservation is granted.

Reservation queuing is not supported by the drive. If a LUN is reserved and a Reserve command is issued from a different initiator, the Target responds with a reservation Conflict.

# 7.30 REZERO UNIT (01)

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | **7** | **6** | **5** | **4** | **3** | **2** | **1** | **0** |
| 0 | Command Code = 01h | | | | | | | |
| 1 | Reserved = 0 | | | Reserved = 0 | | | | |
| 2<br>3<br>4 | Reserved = 0 | | | | | | | |
| 5 | VU = 0 | | Reserved = 0 | | | | FLAG | LINK |

Figure 149. REZERO UNIT (01)

The REZERO UNIT command requests that the target seek logical block address 0.

# 7.31 SEEK (6) - (0B)

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | **7** | **6** | **5** | **4** | **3** | **2** | **1** | **0** |
| 0 | Command Code = 0Bh | | | | | | | |
| 1 | Reserved = 0 | | | (MSB) LBA | | | | |
| 2<br>3 | Logical Block Address (LSB) | | | | | | | |
| 4 | Reserved = 0 | | | | | | | |
| 5 | VU = 0 | | Reserved = 0 | | | | FLAG | LINK |

Figure 150. SEEK (6) - (0B)

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

# 7.32 SEEK (10) - (2B)

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | **7** | **6** | **5** | **4** | **3** | **2** | **1** | **0** |
| 0 | Command Code = 2Bh | | | | | | | |
| 1 | Reserved = 0 | | | Reserved = 0 | | | | 0 |
| 2<br>3<br>4<br>5 | (MSB)<br>Logical Block Address<br>(LSB) | | | | | | | |
| 6<br>7<br>8 | Reserved = 0 | | | | | | | |
| 9 | VU = 0 | | Reserved = 0 | | | | FLAG | LINK |

Figure 151. SEEK (10) - (2B)

The SEEK (10) command requests the drive to seek the specified logical block address.

# 7.33 SEND DIAGNOSTIC (1D)

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
| 0 | Command Code = 1Dh | | | | | | | |
| 1 | Function Code | | | PF | RSVD=0 | SlfTst | Dev0fl | Unt0fl |
| 2 | Reserved = 0 | | | | | | | |
| 3<br>4 | (MSB) Parameter List Length (LSB) | | | | | | | |
| 5 | VU = 0 | | Reserved = 0 | | | | FLAG | LINK |

Figure 152. SEND DIAGNOSTIC (1D)

The SEND DIAGNOSTIC command requests the drive to perform its self-diagnostic test or to perform a function based on a page of information sent in a Data Out phase during the command.

- **PF (Page Format)** bit set to one indicates the data sent by the Initiator conform to the page structure as specified in SCSI-2 standard. This bit must be set to one if the SlfTst bit is set to zero. This bit is ignored by the Target if the SlfTst bit is set.
- **SlfTst** set to one indicates that the device performs its default self-test. If SlfTst is one, the Function code field is ignored. If SlfTst is set to zero, the action to perform is specified in Function code field.

| Value | Function Name | Description |
|---|---|---|
| 000b | NA | Value to be used when the slfTst bit is set to one or if The SEND DIAGNOSTIC command is not invoking one of the other self-test function codes. |
| 001b | Background Short self-test | The device server starts its short self-test routine in background mode. |
| 010b | Background extended self-test | The device server starts its extended self-test routine in background mode. |
| 011b | NA | Reserved |
| 100b | Abort background self-test | Abort the current self-test in the background mode. This value is only valid if a previous SEND DIAGNOSTIC command specified a background self-test function and that function has not been com-pleted. |
| 101b | Foreground short self-test | The device server starts its short self-test routine in the foreground mode. |
| 110b | Foreground extended self-test | The device server starts its extended self-test routine in the foreground mode. |
| 111b | | Reserved |

Figure 153. SEND DIAGNOSTIC Function Code (1D)

- **DevOfl** is ignored by the target for compatibility.
- **UntOfl** is ignored by the target for compatibility.
- **Parameter List Length** must be 0 when the SlfTst bit is one. Otherwise, CHECK CONDITION status will be generated with a sense key of ILLEGAL REQUEST and additional sense of INVALID FIELD IN CDB. If the SlfTst bit is zero, it should be set to the length of the page to be transferred in the DATA OUT phase of the command. If it does not match the expected length of the page a CHECK CONDITION status will be also generated with a sense key of ILLEGAL REQUEST and additional sense of INVALID FIELD IN CDB.

If the motor is not running at the correct speed when the command is received, it is rejected by a CHECK CONDITION status with a NOT READY sense key.

If a fault is detected during the default or foreground self-test, a CHECK CONDITION is reported as an end status. If a fault is detected during the background self-test, it is logged in the log page for later retrieval by a LOG SENSE command.

See Section 10.17, "Diagnostics" on page 215 for detailed listing of operations carried out by SEND DIAGNOSTIC command and Power on Diagnostics.

## 7.33.1 Send Diagnostic Page 0

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

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

Figure 154. Diagnostic Page 0

## 7.33.2 Send Diagnostic Page 40

This allows the initiator to translate a logical block address or physical sector address to the other format. The address to be translated is passed to the target with the Send Diagnostic command and the results are returned to the initiator by the Receive Diagnostics command.

The target will read the parameter list from the initiator and, if no errors are detected in the parameter list, Good Status will be returned. The data translation will be performed upon receipt of the Receive Diagnostics command.

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
| 0 | Page Code = 40h | | | | | | | |
| 1 | Reserved = 0 | | | | | | | |
| 2-3 | Page Length = 0Ah | | | | | | | |
| 4 | Reserved = 0 | | | | | Supplied format | | |
| 5 | Reserved = 0 | | | | | Translate format | | |
| 6-13 | Address to Translate | | | | | | | |

Figure 155. Diagnostic Page 40

**Supplied Format** may take one of the following three values:

- **000b** Block format
- **100b** Bytes From Index format
- **101b** Physical Sector format

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

- **Translate Format** specifies the format that the address should be translated into. If the supplied format is the Block format, the Translate format must be either Bytes from Index or Physical Sector format. If the supplied format is the Bytes from Index or Physical Sector format, the Translate format must be Block format. If either of the format fields is invalid or they specify the same format, the

command will terminate with CHECK CONDITION status with a Sense of Illegal Request and Illegal Field in Parameter List.

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

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
| 6-8 | Cylinder Number | | | | | | | |
| 9 | Head Number | | | | | | | |
| 10-13 | Sector Number or Bytes from Index | | | | | | | |

Figure 156. Address to translate

## 7.33.3 Send Diagnostic Page 80

This requests the drive to perform off-line read scan.

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
| 0 | Page Code = 80h | | | | | | | |
| 1 | Reserved = 0 | | | | | | | |
| 2-3 | Page Length = 04h | | | | | | | |
| 4 | SMART Revision = 03h | | | | | | | |
| 5 | Reserved = 0 | | | | | | | |
| 6<br>7 | (MSB) Off-line Immediate Time (LSB) | | | | | | | |

Figure 157. Send Diagnostic Page 80

- **Off-line Immediate Time** specifies the delay to start off-line read scan in seconds. When this value is zero, the target starts off-line read scan immediately.

# 7.34 SET DEVICE IDENTIFIER (A4)

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | **7** | **6** | **5** | **4** | **3** | **2** | **1** | **0** |
| 0 | Command Code = A4h | | | | | | | |
| 1 | Reserved = 0 | | | | Service Action = 06h | | | |
| 2<br>3<br>4<br>5 | Reserved = 0 | | | | | | | |
| 6<br>\|<br>9 | (MSB)<br>Parameter List Length<br>(LSB) | | | | | | | |
| 10 | Reserved = 0 | | | | | | | |
| 11 | VU = 0 | | Reserved = 0 | | | | FLAG | LINK |

Figure 158. SET DEVICE IDENTIFIER(A4)

The SET DEVICE IDENTIFIER command requests that the device identifier information be set to the value received in the SET DEVICE IDENTIFIER parameter list.

On successful completion of a SET DEVICE IDENTIFIER command a unit attention is generated for all initiators except the one that issued the service action. When reporting the unit attention condition the additional sense code is set to DEVICE IDENTIFIER CHANGED.

- **Parameter List Length** field specifies the length in bytes of the Identifier that is transferred from the host system to the target. The maximum value for this field is 512 bytes. A parameter list length of zero indicates that no data is transferred, and that subsequent REPORT DEVICE IDENTIFIER commands return an Identifier length of zero.

The SET DEVICE IDENTIFIER parameter list contains the identifier to be set by the addressed logical unit.

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | **7** | **6** | **5** | **4** | **3** | **2** | **1** | **0** |
| 0<br>n | Identifier | | | | | | | |

Figure 159. SET DEVICE IDENTIFIER Parameter List

The IDENTIFIER field is a vendor specific value, to be returned in subsequent REPORT DEVICE IDENTIFIER commands.

# 7.35 START STOP UNIT (1B)

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | **7** | **6** | **5** | **4** | **3** | **2** | **1** | **0** |
| 0 | Command Code = 1Bh | | | | | | | |
| 1 | Reserved = 0 | | | Reserved = 0 | | | | Immed |
| 2<br>3 | Reserved = 0 | | | | | | | |
| 4 | Power Conditions = 0 | | | | Reserved=0 | | LoEj = 0 | Start |
| 5 | VU = 0 | | Reserved = 0 | | | | FLAG | LINK |

Figure 160. START STOP UNIT (1B)

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

- **Immed** bit is to specify

  **0** status is to be returned at the end of the operation
  **1** GOOD status shall always be returned immediately after command has been received. The TEST UNIT READY command may be used to determine when the drive becomes ready after a spin-up.

- **Power Conditions** is not supported by the drive and must be set to 0.
- **LoEj** is not supported by the drive and must be set to 0.
- **Start** bit is to specify:

  **0** stop the spindle
  **1** start the spindle

**Note:** Once the drive has become ready (after a power on), the Start stop unit command can be used without any errors regardless of the state of the motor: stopped or spinning.

# 7.36 SYNCHRONIZE CACHE (35)

| Byte | BIT | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | **7** | **6** | **5** | **4** | **3** | **2** | **1** | **0** |
| 0 | Command Code = 35h | | | | | | | |
| 1 | Reserved = 0 | | | Reserved = 0 | | | Immed = 0 | RelAdr = 0 |
| 2<br>3<br>4<br>5 | (MSB)<br>Logical Block Address<br>(LSB) | | | | | | | |
| 6 | Reserved = 0 | | | | | | | |
| 7<br>8 | (MSB)<br>Number of Blocks<br>(LSB) | | | | | | | |
| 9 | VU = 0 | | Reserved = 0 | | | | FLAG | LINK |

Figure 161. SYNCHRONIZE CACHE (35)

The SYNCHRONIZE CACHE Command ensures that logical blocks in the cache have their most recent data value recorded on the media.

- **Logical Block Address** is to specify where the operation is to begin.
- **Number of Blocks** specifies the total number of contiguous logical blocks within the range. Number of Blocks of zero indicates that all remaining logical blocks on the logical unit shall be within the range.
- **Immed** (immediate) must be zero. An immediate bit of zero indicates that the status shall not be returned until the operation has completed. If the Immed bit is set to one, the drive returns a Check Condition status. The sense key shall be set to Illegal Request and the additional sense code shall be set to Invalid Field in CDB.
- **RelAdr** (Relative Address) must be zero. The drive does not support the relative addressing. If the RelAdr bit is set to one, the drive returns Check Condition status. The sense key shall be set to Illegal Request and the additional sense code shall be set to Invalid Field in CDB.

# 7.37 TEST UNIT READY (00)

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | **7** | **6** | **5** | **4** | **3** | **2** | **1** | **0** |
| 0 | Command Code = 00h | | | | | | | |
| 1 | Reserved = 0 | | | Reserved = 0 | | | | |
| 2<br>3<br>4 | Reserved = 0 | | | | | | | |
| 5 | VU = 0 | | Reserved = 0 | | | | FLAG | LINK |

Figure 162. TEST UNIT READY (00)

The TEST UNIT READY command allows the initiator to check if the drive is READY. The SCSI specification defines READY as the condition where the device will accept a media-access command without returning CHECK CONDITION status.

The drive will first verify that the motor is spinning at the correct speed. If the spindle motor is not spinning at the correct speed, CHECK CONDITION status is returned with sense key of NOT READY. If the motor is spinning at the correct speed, the drive accepts normal media access commands.

The TEST UNIT READY command is not intended as a diagnostic. No self diagnostic is performed by the device as a result of this command.

The TEST UNIT READY command has special significance for power sequencing using the UNIT START command with an Immediate bit of one. In this mode the UNIT START command returns TASK COMPLETE status before the completion of motor spin-up and expects the initiator to issue TEST UNIT READY commands to determine when the motor has reached the proper speed.

**Note:** The spindle automatically starts in automatic spin-up Mode. The drive does not execute any commands other than TEST UNIT READY, INQUIRY, or REQUEST SENSE command until the Power On sequence is complete. The drive will return CHECK CONDITION status with NOT READY sense key and IN PROCESS OF BECOMING READY sense code for all other commands during the Power On period.

# 7.38 VERIFY (2F)

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | **7** | **6** | **5** | **4** | **3** | **2** | **1** | **0** |
| 0 | Command Code = 2Fh | | | | | | | |
| 1 | Reserved = 0 | | | DPO | Reserved = 0 | | Byte Chk | RSVD =0 |
| 2<br>3<br>4<br>5 | (MSB) Logical Block Address (LSB) | | | | | | | |
| 6 | Reserved = 0 | | | | | | | |
| 7<br>8 | (MSB) Transfer Length (LSB) | | | | | | | |
| 9 | VU = 0 | | Reserved = 0 | | | | FLAG | LINK |

Figure 163. VERIFY (2F)

The VERIFY command requests that the drive verify the data written on the media. A verification length of zero indicates that no data will be transferred. This condition is not considered an error.

- **ByteChk** bit set to zero indicates that the data is read from the disk and verified using ECC.If an ECC error is detected in the verify process, CHECK CONDITION status is returned with sense key set to MEDIUM ERROR. ByteChk bit set to one indicates that byte-by-byte comparison is performed between the data on the disk and data transferred from the initiator during the data-out phase.

  If the comparison is unsuccessful, the command is terminated with CHECK CONDITION status and the sense key is set to MISCOMPARE.

**DPO** A DPO (Disable Page Out) bit of one indicates that the data accessed by this command is to be assigned the lowest priority for being written into or retained by the cache. A DPO bit of one overrides any retention priority specified in the Mode Select Page 8 Caching Parameters. A DPO bit of zero indicates the priority is determined by the retention priority. The initiator should set the DPO bit when the blocks read by this command are not likely to be read again in the near future.

# 7.39 WRITE (6) - (0A)

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | **7** | **6** | **5** | **4** | **3** | **2** | **1** | **0** |
| 0 | Command Code = 0Ah | | | | | | | |
| 1 | Reserved = 0 | | | (MSB) LBA | | | | |
| 2<br>3 | Logical Block Address<br>(LSB) | | | | | | | |
| 4 | Transfer Length | | | | | | | |
| 5 | VU = 0 | | Reserved = 0 | | | | FLAG | LINK |

Figure 164. WRITE (6) - (0A)

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

See **READ (6) - (08)** on page 127 for the parameters.

# 7.40 WRITE (10) - (2A)

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | **7** | **6** | **5** | **4** | **3** | **2** | **1** | **0** |
| 0 | Command Code = 2Ah | | | | | | | |
| 1 | Reserved = 0 | | | DPO | FUA | Reserved=0 | | RelAdr = 0 |
| 2<br>3<br>4<br>5 | (MSB)<br>Logical Block Address<br>(LSB) | | | | | | | |
| 6 | Reserved = 0 | | | | | | | |
| 7<br>8 | (MSB<br>Transfer Length<br>(LSB) | | | | | | | |
| 9 | VU = 0 | | Reserved = 0 | | | | FLAG | LINK |

Figure 165. WRITE (10) - (2A)

The WRITE (10) command requests that the drive write the data transferred from the initiator. This command is processed like the standard WRITE command except for the longer transfer length.

**Transfer length** The number of contiguous blocks to be transferred. If the transfer length is zero, the seek occurs, but no data is transferred. This condition is not considered an error.

**DPO** A DPO (Disable Page Out) bit of one indicates that the data accessed by this command is to be assigned the lowest priority for being written into or retained by the cache. A DPO bit of one overrides any retention priority specified in the Mode Select Page 8 Caching Parameters. A DPO bit of zero indicates that the priority is determined by the retention priority. The Initiator should set the DPO bit when the blocks written by this command are not likely to be read in the near future.

**FUA** Force unit access. A FUA bit of one indicates that the Target must write the data to the media before returning Good Status. A FUA bit of zero indicates that the Target may return Good Status prior to writing the data to the media.

**RelAdr** Relative Block Address is not supported. Must be set to zero.

# 7.41 WRITE AND VERIFY (2E)

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
| 0 | Command Code = 2Eh | | | | | | | |
| 1 | Reserved = 0 | | | DPO | Reserved = 0 | | Byte Chk | RelAdr = 0 |
| 2<br>3<br>4<br>5 | (MSB)<br>Logical Block Address<br>(LSB) | | | | | | | |
| 6 | Reserved = 0 | | | | | | | |
| 7<br>8 | (MSB)<br>Transfer Length<br>(LSB) | | | | | | | |
| 9 | VU = 0 | | Reserved = 0 | | | | FLAG | LINK |

Figure 166. WRITE AND VERIFY (2E)

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

- **Transfer Length** of zero indicates that no data is transferred.

  If caching is enabled, the command performs an implied Force Unit Access (FUA) and an implied Synchronize Cache before starting the operation. This insures that the medium, not the cache, is being verified.

- **ByteChk** bit set to zero indicates that the data is read back from the disk and verified using ECC after the successful write operation. If an ECC error is detected in the verify process, CHECK CONDITION status is returned with sense key set to MEDIUM ERROR. ByteChk bit set to one indicates that byte-by-byte comparison is performed between data on the disk starting the block specified in LBA field and data transferred from the initiator.

  If the comparison is unsuccessful, the command is terminated with CHECK CONDITION status and the sense key is set to MISCOMPARE.

- **DPO** (Disable Page Out) bit of one indicates that the data written by this command is to be assigned the lowest priority for being written into or retained by the cache. A DPO bit of one overrides any retention priority specified in the Mode Select Page 8 Caching parameters. A DPO bit of zero indicates the priority is determined by the retention priority.

  The initiator should set the DPO bit when the blocks written by this command are not likely to be read again in the near future.

- **Relative Block Address** is not supported. Must be set to zero.

# 7.42 WRITE BUFFER (3B)

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | **7** | **6** | **5** | **4** | **3** | **2** | **1** | **0** |
| 0 | Command Code = 3Bh | | | | | | | |
| 1 | Reserved = 0 | | | | Mode | | | |
| 2 | Buffer ID | | | | | | | |
| 3<br>4<br>5 | (MSB)<br>Buffer Offset<br>(LSB) | | | | | | | |
| 6<br>7<br>8 | (MSB)<br>Parameter List Length<br>(LSB) | | | | | | | |
| 9 | VU = 0 | | Reserved = 0 | | | | FLAG | LINK |

Figure 167. WRITE BUFFER (3B)

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

This command will cause the entire cache to be emptied.

The function of this command and the meaning of fields within the command descriptor block depend on the contents of the mode field.

| MODE | Description |
|---|---|
| **0000** | Write combined header and data |
| **0010** | Data |
| **0100** | Download Microcode |
| **0101** | Download Microcode and Save |
| **0111** | Download Microcode and Save(Single  Binary Chunked) |
| **1010** | Write Data to Echo Buffer |

No other modes are supported by the drive.

## 7.42.1 Combined Header And Data (Mode 0000b)

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

**Buffer ID** must be zero. If another value is specified, no download function is performed and the command is terminated with CHECK CONDITION status. And the drive shall set the sense key to ILLEGAL REQUEST and additional sense code to ILLEGAL FIELD IN CDB.

**Buffer Offset** must be zero. If another value is specified, no download function is performed and the command is terminated with CHECK CONDITION status. And the drive shall set the sense key to ILLEGAL REQUEST and additional sense code to ILLEGAL FIELD IN CDB.

**Parameter List Length** specifies the number of bytes that shall be transferred during the DATA OUT phase. This number includes four bytes of header, so the data length to be stored in the drive buffer is transfer length minus four. If the length exceeds the buffer size, the command is terminated with CHECK CONDITION status. And the drive shall set sense key to ILLEGAL REQUEST and additional sense code to ILLEGAL FIELD IN CDB.

The four-byte header consists of all reserved bytes.

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
| 0 | Reserved = 0 | | | | | | | |
| 1 | Reserved = 0 | | | | | | | |
| 2 | Reserved = 0 | | | | | | | |
| 3 | Reserved = 0 | | | | | | | |

Figure 168. WRITE BUFFER Header

## 7.42.2 Write Data (Mode 0010b)

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

**Buffer ID** must be set to zero, indicating the data transfer buffer. If another value is specified, the command is terminated with CHECK CONDITION status. The drive shall set the sense key to ILLEGAL REQUEST and additional the sense code to ILLEGAL FIELD IN CDB.

**Buffer Offset** specifies the offset of the memory space specified by the Buffer ID. The initiator should conform to the offset boundary requirements returned in the READ BUFFER descriptor. If the value exceeds the buffer specified, the command is terminated with CHECK CONDITION status. The drive shall set the sense key to ILLEGAL REQUEST and additional sense code to ILLEGAL FIELD IN CDB.

**Parameter List Length** specifies the Parameter List Length. It must be

- less than the capacity of the buffer size and
- on a sector boundary.

If an invalid value is specified, the command is terminated with CHECK CONDITION status. The drive shall set the sense key to ILLEGAL REQUEST and additional sense code to ILLEGAL FIELD IN CDB.

## 7.42.3 Download Microcode (Mode 0100b)

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

**Buffer ID** field is used to indicate which portion of the microcode image is being downloaded. The following Buffer IDs are supported by the Target:

- 00h : Main Microprocessor Code
- 81h : Reserved Area Data

Any other value for the Buffer ID will cause the command to terminate with CHECK CONDITION status. The drive shall set the sense key to ILLEGAL REQUEST and additional sense code to ILLEGAL FIELD IN CDB.

**Buffer Offset** must be zero. If an invalid value is specified, the command is terminated with CHECK CONDITION status. The drive shall set the sense key to ILLEGAL REQUEST and additional sense code to ILLEGAL FIELD IN CDB.

**Parameter List Length** must be the size of the data set to be downloaded. It may also be set to 0000h in which case no code is updated. If an invalid value is specified, the command is terminated with CHECK CONDITION status. The drive shall set the sense key to ILLEGAL REQUEST and additional sense code to ILLEGAL FIELD IN CDB.

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

**Note:** The Download Microcode mode described in this specification is to indicate that the drive will accept a command with this mode, though it is not expected that a user will ever issue such a command. To use the write buffer command with this mode, a special microcode version is required from development. If such a microcode is released from development, then it will include appropriate instructions on the function of new microcode and its effect on the drive operations after download.

**Note:** If the write buffer command with this mode is executed and the invalid code is downloaded, then it is to be expected that the drive will never be complete or hang up the subsequent command. This condition is normally recoverable by a power on/off cycle, but there is no guarantee of it.

## 7.42.4 Download Microcode and Save (Mode 0101b)

In this mode the data is transferred to the drive to save into the System reserved area on the disk. This is for functional upgrade and configuration change reflecting the user's requirements and the manufacturer's reason or both, and it is stored in the media as a permanent copy. The newly downloaded code becomes effective after the drive issues and completes a self-initiated Power On Reset.

**Note:** It requires up to 30 seconds to update the microcode including the Flash ROM update.

**Note:** New code to be downloaded to the drive will be provided by development either by request of a customer for an additional function or as a result of a functional change by development. However please note that not all possible fixes or new functions can be applied to a drive in this manner and that there is a very high dependency on the level of ROM code contained within the drive. If an invalid code or a code not compatible with the ROM code is downloaded, the drive will usually reject this code and will continue normal operation. However there is a small possibility that an invalid code will be accepted. If this occurs, the unit usually becomes inoperable and will have to be returned to the manufacturer for recovery.

**Buffer ID** field is used to indicate which portion of the microcode image is being downloaded. The following Buffer IDs are supported by the Target:

- 00h : Main Microprocessor Code with all others in one (Single Binary or Chunked)
- 01h - 02h : Reserved
- 80h - 82h : Reserved

Any other value for the Buffer ID will cause the command to terminate with CHECK CONDITION status. The drive shall set the sense key to ILLEGAL REQUEST and additional sense code to ILLEGAL FIELD IN CDB.

### 7.42.4.1 Download Microprocessor Microcode and Save (Buffer ID = 00h)

Two types of download function are supported: Single Binary Download and Single Binary Chunked.

**Buffer ID** is 00h. (Single Binary Download)
The first one is that the data set consist of all necessary code and tables in a one big data set. This is starting with the special header data followed by tables, RAM code, three Overlay codes, EEPROM data, and Flash-ROM code. (Flash-ROM code is optional depends on the card type and version if card has update capability.) This download function is available when the drive is running with full function (normal running) before this write buffer is issued.

It requires up to 30 seconds to finish the command. After completion of this type of write buffer command the drive will start as "Power on Reset" and running with newly downloaded code and configuration. Therefore the initiator may need special treatment for this target.

**Buffer ID** is 00h. (Single Binary Chunked)
The second uses the same data set as the first (Single Binary). It separates one big binary file into segments of 32K bytes and then issues a Write buffer command with ID=00 repeatedly with the same number as the segment. The last segment may be equal to or less than 32K bytes. (For example, if the single binary file size is 426K bytes, there are 13 files of 32K bytes. Thus 13 continuous Write Buffer commands with ID=00 should be issued.) After issuing continuous Write Buffer for all segments, the drive behavior is the same as the single binary.

## 7.42.5 Download Microcode and Save (Mode 0111b)

**Buffer Offset** is 00h. (Single Binary Chunked Mode 7)
The 3rd uses the same data set as the first (Single Binary). It separates one big binary file into segments of 32K bytes and then issues a Write buffer command with Offset=00 repeatedly with the same number as the segment. The last segment may be equal to or less than 32K bytes. (For example, if the single binary drive size is 426K bytes, there are thirteen files of 32K bytes. Thus thirteen continuous Write Buffer commands with Offset=00 should be issued.) After issuing continuous Write Buffer for all segments, the drive behavior is the same as the single binary.

## 7.42.6 Write Data to Echo Buffer (Mode 1010b)

In this mode the target transfers data into the echo buffer. The echo buffer is assigned in the same manner by the target as it would for a write operation. Data will be sent aligned on four-byte boundaries.

Upon successful completion of a Write Buffer command the data will be preserved in the echo buffer unless there is an intervening command to any logical unit in which case it may be changed.

# 7.43 WRITE LONG (3F)

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
| 0 | Command Code = 3Fh | | | | | | | |
| 1 | Reserved = 0 | | | Reserved = 0 | | | | RelAdr = 0 |
| 2<br>3<br>4<br>5 | (MSB)<br>Logical Block Address<br>(LSB) | | | | | | | |
| 6 | Reserved = 0 | | | | | | | |
| 7<br>8 | (MSB) Byte Transfer Length<br>(LSB) | | | | | | | |
| 9 | VU = 0 | | Reserved = 0 | | | | FLAG | LINK |

Figure 169. WRITE LONG (3F)

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

- User Data
- 40 bytes of ECC data

Parameters are

- **Logical Block Address** field specifies the logical block at which the write operation shall occur.
- **Byte Transfer Length**. This field must specify the exact number of bytes of data available for transfer. If a non-zero byte transfer length does not match the available data length, the target terminates the command with CHECK CONDITION status, then the sense key is set to ILLEGAL REQUEST, and an additional sense code is set to INVALID FIELD IN CDB. The valid and ILI bits are set to one and the information field is set to the difference of the requested length minus the actual length in bytes. Negative values are indicated by two's complement notation.
- **RelAdr** (Relative Block Address) is not supported. Must be set to ZERO.

# 7.44 WRITE SAME (41)

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | **7** | **6** | **5** | **4** | **3** | **2** | **1** | **0** |
| 0 | Command Code = 41h | | | | | | | |
| 1 | Reserved = 0 | | | Reserved = 0 | | | | RelAdr = 0 |
| 2<br>3<br>4<br>5 | (MSB) Logical Block Address (LSB) | | | | | | | |
| 6 | Reserved = 0 | | | | | | | |
| 7<br>8 | (MSB) Number of Blocks (LSB) | | | | | | | |
| 9 | VU = 0 | | Reserved = 0 | | | | FLAG | LINK |

Figure 170. WRITE SAME (41)

The Write Same command instructs the Target to write a single block of data transferred to the Target from the Initiator to a number of sequential logical blocks. This command is useful for writing large data areas without sending all of the data over the SCSI bus.

- **Logical Block Address** specifies the address at which the write begins. The Number of Blocks specifies the number of contiguous blocks to be written. If the number is zero, all of the remaining blocks on the specified Logical Unit are written.
- **RelAdr (Relative Block Address)** is not supported and must be set to be ZERO.

The data for this command is not retained in the cache.

# 8.0 SCSI Status Byte

Upon the completion of a command a status byte is sent to the initiator. Additional sense information may also be available depending on the contents of the status byte. The following section describes the possible values for the status byte and sense data. All Reserved fields (R) are set to zero.

| Bit | | | | | | | |
|---|---|---|---|---|---|---|---|
| 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
| Reserved = 0 | | Status Code | | | | | RSVD |

Figure 171. SCSI Status Byte. Format of the SCSI STATUS byte.

| STATUS BYTE | Description |
|---|---|
| **00h** | GOOD<br>The command has been successfully completed. |
| **02h** | CHECK CONDITION<br>An error, exception, or abnormal condition has been detected. The sense data is set by the drive. The REQUEST SENSE command should be issued to determine the nature of the condition. |
| **04h** | CONDITION MET<br>This status is returned when an unlinked PRE-FETCH command has been successfully completed. |
| **08h** | BUSY<br>This condition is returned when disconnect privilege is not granted while the drive is BUSY processing the other command for the other initiator. The normal initiator recovery action is to issue the command at a later time or to reissue the command and grant the disconnect privilege. |
| **10h** | INTERMEDIATE<br>This status is returned for every command except PRE-FETCH command in a series of linked commands (except the last command), unless the command is terminated with CHECK CONDITION, RESERVATION CONFLICT, BUSY status. If INTERMEDIATE or INTERMEDIATE-CONDITION MET status is not returned, the series of linked commands is terminated and the task is ended. |
| **14h** | INTERMEDIATE CONDITION MET<br>This status is returned when a linked PRE-FETCH command has been completed, unless the command is terminated with CHECK CONDITION, RESERVATION CONFLICT, BUSY status. If INTERMEDIATE or INTERMEDIATE-CONDITION MET status is not returned, the series of linked commands is terminated and the task is ended. |
| **18h** | RESERVATION CONFLICT<br>This status is returned whenever an SCSI device attempts to access the drive, but it has been reserved by another initiator. (See <Command.lwp\Interface:Reserve> on page 0.) |

**28h** QUEUE FULL

This status indicates that the target's command queue is full. If a tagged command queuing feature is enabled and there is no room on the command queue, this status is returned when the initiator sends a command. For this status sense is not valid.

# 9.0 SCSI Message System

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

## 9.1 Supported Messages

The messages supported by the drive are listed in the figure below.

| MESSAGE | CODE(hex) | Direction | | Negate ATN Before last ACK |
|---|---|---|---|---|
| TASK COMPLETE | 00 | IN | | --- |
| SYNCHRONOUS DATA TRANSFER REQUEST | 010301 | IN | OUT | Yes |
| WIDE DATA TRANSFER REQUEST* | 010203 | IN | OUT | Yes |
| PARALLEL PROTOCOL REQUEST | 010604 | IN | OUT | Yes |
| SAVE DATA POINTER | 02 | IN | | --- |
| RESTORE POINTERS | 03 | IN | | --- |
| DISCONNECT | 04 | IN | | --- |
| INITIATOR DETECTED ERROR | 05 | | OUT | Yes |
| ABORT TASK SET | 06 | | OUT | Yes |
| MESSAGE REJECT | 07 | IN | OUT | Yes |
| NO OPERATION | 08 | | OUT | Yes |
| MESSAGE PARITY ERROR | 09 | | OUT | Yes |
| LINKED TASK COMPLETE | 0A | IN | | --- |
| LINKED TASK COMPLETE (w/FLAG) | 0B | IN | | --- |
| TARGET RESET | 0C | | OUT | Yes |
| ABORT TASK | 0D | | OUT | Yes |
| CLEAR TASK SET | 0E | | OUT | Yes |
| SIMPLE | 20XX | IN | OUT | No |
| HEAD OF QUEUE | 21XX | | OUT | No |
| ORDERED | 22XX | | OUT | No |
| IGNORE WIDE RESIDUE* | 2301 | IN | | --- |
| IDENTIFY | 80-FF | IN | | --- |
| IDENTIFY | 80-FF | | OUT | No |

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

Figure 172. Supported Messages

If an unsupported message is received, the drive will send the MESSAGE REJECT message to the initiator. If at the time the unsupported message is received a valid NEXUS exists, the drive will continue with the command. If no valid NEXUS exists, the drive will go to Bus Free.

## 9.1.1 TASK COMPLETE (00)

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

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

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

Figure 173. Synchronous Data Transfer Request.

A pair of Synchronous Data Transfer Request (SDTR) messages shown in Figure 159 are exchanged between an Initiator and a Target to establish the synchronous data transfer mode between the two devices. The message exchange establishes the permissible transfer period and REQ/ACK offset for a synchronous data transfer between the two devices. The initiator may initiate a synchronous data transfer negotiation at any time after the LUN has been identified. A Synchronous Data Transfer Request (SDTR) message exchange shall be initiated by a SCSI device whenever a previously arranged data transfer agreement may have become invalid. *The agreement becomes invalid after any condition which may leave the data transfer agreement in an indeterminate state such as*

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

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

**M** The transfer period (M) 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 drive supports transfer period in the range of 50 ns to 425 ns in increments of 25 ns. In addition to this, when the drive is working in LVD mode, it supports transfer period of 25 ns.

**REQ/ACK Offset**

**X** The ACK/REQ offset (X above) is the maximum number of REQ pulses allowed to be outstanding before the leading edge of its corresponding ACK pulses is received at the drive. A REQ/ACK offset value of zero indicates asynchronous data transfer mode. The drive supports REQ/ACK offset values in the range 0 through 63.

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

### 9.1.2.1 Synchronous Negotiation started by the Initiator

When target respond with REQ/ACK offset value of 0, the initiator shall use asynchronous data transfer mode.

**LVD mode.** The Target responds to each Initiator requested transfer period as shown below.

| Initiator Request | Target Response | Target Transfer Period (ns) | Maximum Burst Rate (MT/s) |
|---|---|---|---|
| 0 <= Mi <= 09 | Mt = 10 | 25 | 40.00 |
| 10 <= Mi <= 10 | Mt = 10 | 25 | 40.00 |
| 11 <= Mi <= 12 | Mt = 12 | 50 | 20.00 |
| 13 <= Mi <= 18 | Mt = Mi | 75 | 13.33 |
| 19 <= Mi <= 25 | Mt = Mi | 100 | 10.00 |
| 26 <= Mi <= 31 | Mt = Mi | 125 | 8.00 |
| 32 <= Mi <= 37 | Mt = Mi | 150 | 6.67 |
| 38 <= Mi <= 43 | Mt = Mi | 175 | 5.71 |
| 44 <= Mi <= 50 | Mt = Mi | 200 | 5.00 |
| 51 <= Mi <= 56 | Mt = Mi | 225 | 4.44 |
| 57 <= Mi <= 62 | Mt = Mi | 250 | 4.00 |
| 63 <= Mi <= 68 | Mt = Mi | 275 | 3.64 |
| 69 <= Mi <= 75 | Mt = Mi | 300 | 3.33 |
| 76 <= Mi <= 81 | Mt = Mi | 325 | 3.08 |
| 82 <= Mi <= 87 | Mt = Mi | 350 | 2.86 |
| 88 <= Mi <= 93 | Mt = Mi | 375 | 2.67 |
| 94 <= Mi <= 100 | Mt = Mi | 400 | 2.50 |
| 101 <= Mi <= 106 | Mt = Mi | 425 | 2.35 |
| 107 <= Mi <= 255 | Mt = 106 | (Asynch mode) | N/A |

Figure 174. Initiator Request/Target Response (LVD mode)

**SE mode.** The Target responds to each Initiator requested transfer period as shown below.

| Initiator Request | Target Response | Target Transfer Period (ns) | Maximum Burst Rate (MT/s) |
|---|---|---|---|
| 0 <= Mi <= 11 | Mt = 12 | 50 | 20.00 |
| 12 <= Mi <= 12 | Mt = 12 | 50 | 20.00 |
| 13 <= Mi <= 18 | Mt = Mi | 75 | 13.33 |
| 19 <= Mi <= 25 | Mt = Mi | 100 | 10.00 |
| 26 <= Mi <= 31 | Mt = Mi | 125 | 8.00 |
| 32 <= Mi <= 37 | Mt = Mi | 150 | 6.67 |
| 38 <= Mi <= 43 | Mt = Mi | 175 | 5.71 |
| 44 <= Mi <= 50 | Mt = Mi | 200 | 5.00 |
| 51 <= Mi <= 56 | Mt = Mi | 225 | 4.44 |
| 57 <= Mi <= 62 | Mt = Mi | 250 | 4.00 |
| 63 <= Mi <= 68 | Mt = Mi | 275 | 3.64 |
| 69 <= Mi <= 75 | Mt = Mi | 300 | 3.33 |
| 76 <= Mi <= 81 | Mt = Mi | 325 | 3.08 |
| 82 <= Mi <= 87 | Mt = Mi | 350 | 2.86 |
| 88 <= Mi <= 93 | Mt = Mi | 375 | 2.67 |
| 94 <= Mi <= 100 | Mt = Mi | 400 | 2.50 |
| 101 <= Mi <= 106 | Mt = Mi | 425 | 2.35 |
| 107 <= Mi <= 255 | Mt = 106 | (Asynch mode) | N/A |

Figure 175. Initiator Request/Target Response (SE mode)

### 9.1.2.2 Synchronous Negotiation started by the Target

If the drive recognizes that negotiation is required, and if the "Target Initiated SDTR" jumper is installed, the drive sends a SDTR message to the initiator with minimum transfer period on the current receiver mode. The drive interprets the Initiator corresponding transfer period as shown in the figure below.

**LVD mode**

| Initiator's Request | Target Transfer Period (ns) | Maximum Burst Rate (MT/s) |
|---|---|---|
| 0 <= Mi <= 09 | Send M+=10 to negotiate | N/A |
| 10 <= Mi <= 10 | 25 | 40.00 |
| 11 <= Mi <= 11 | 50 | 20.00 |
| 12 <= Mi <= 12 | 50 | 20.00 |
| 13 <= Mi <= 18 | 75 | 13.33 |
| 19 <= Mi <= 25 | 100 | 10.00 |
| 26 <= Mi <= 31 | 125 | 8.00 |
| 32 <= Mi <= 37 | 150 | 6.67 |
| 38 <= Mi <= 43 | 175 | 5.71 |
| 44 <= Mi <= 50 | 200 | 5.00 |
| 51 <= Mi <= 56 | 225 | 4.44 |
| 57 <= Mi <= 62 | 250 | 4.00 |
| 63 <= Mi <= 68 | 275 | 3.64 |
| 69 <= Mi <= 75 | 300 | 3.33 |
| 76 <= Mi <= 81 | 325 | 3.08 |
| 82 <= Mi <= 87 | 350 | 2.86 |
| 88 <= Mi <= 93 | 375 | 2.67 |
| 94 <= Mi <= 100 | 400 | 2.50 |
| 101 <= Mi <= 106 | 425 | 2.35 |
| 107 <= Mi <= 255 | Send Message Reject(Asynch mode) | N/A |

Figure 176. Target Response to Initiator's Transfer Period (LVD mode)

**SE Mode**

| Initiator's Request | Target Transfer Period (ns) | Maximum Burst Rate (MT/s) |
|---|---|---|
| 0 <= Mi <= 11 | M+=12 to negotiate | N/A |
| 12 <= Mi <= 12 | 50 | 20.00 |
| 13 <= Mi <= 18 | 75 | 13.33 |
| 19 <= Mi <= 25 | 100 | 10.00 |
| 26 <= Mi <= 31 | 125 | 8.00 |
| 32 <= Mi <= 37 | 150 | 6.67 |
| 38 <= Mi <= 43 | 175 | 5.71 |
| 44 <= Mi <= 50 | 200 | 5.00 |
| 51 <= Mi <= 56 | 225 | 4.44 |
| 57 <= Mi <= 62 | 250 | 4.00 |
| 63 <= Mi <= 68 | 275 | 3.64 |
| 69 <= Mi <= 75 | 300 | 3.33 |
| 76 <= Mi <= 81 | 325 | 3.08 |
| 82 <= Mi <= 87 | 350 | 2.86 |
| 88 <= Mi <= 93 | 375 | 2.67 |
| 94 <= Mi <= 100 | 400 | 2.50 |
| 101 <= Mi <= 106 | 425 | 2.35 |
| 107 <= Mi <= 255 | Send Message Reject (Asynch mode) | N/A |

Figure 177. Target Response to Initiator's Transfer Period (SE mode)

## 9.1.3 WIDE DATA TRANSFER REQUEST (01,02,03H)

A pair of Wide Data Transfer Request messages is exchanged between an Initiator and a Target to establish a data transfer width agreement between the two devices. The Initiator may initiate a wide data transfer negotiation at any time after the LUN has been identified. The Target initiates a wide data transfer negotiation if the Target has not negotiated with the Initiator since the last time the Target was Reset (Power-on Reset, SCSI Bus Hard Reset, or Target Reset message).

Target-initiated negotiation occurs either immediately following the Command phase or immediately following the first reconnection. In either case negotiation occurs before any Data phase between the Target and the Initiator. The Target will negotiate the data transfer width agreement prior to negotiating the synchronous data transfer agreement. If a synchronous data transfer agreement is in effect when a Wide Data Transfer Request message is received, the Target will reset the synchronous agreement to asynchronous mode.

The implied data transfer width agreement remains in effect until the Target is Reset (Power-on Reset, SCSI Bus "hard" Reset, or Target Reset message) or a new data transfer width agreement is negotiated. If a Reset occurs, the Target goes to eight bit mode.

| Byte | Value | Description |
|---|---|---|
| 0 | 01H | Extended message |
| 1 | 02H | Extended message length |
| 2 | 03H | WIDE DATA TRANSFER REQUEST code |
| 3 | E | Transfer width exponent |

Figure 178. Wide Data Transfer Request.

**E** The Transfer Width Exponent (E) is two to the transfer width exponent bytes wide. Valid data transfer widths are 8 bits (E = 00h) and 16 bits (E = 01h). Value of E greater than 01h are reserved.

### 9.1.3.1 Transfer Width Negotiation started by the Initiator

If the Initiator recognizes that negotiation is required and sends a Wide Data Transfer Request message out, the Target responds by changing to the Message In phase and sending a Wide Data Transfer Request message in to the Initiator prior to transferring any additional message bytes (or any other Information phase bytes) from the Initiator. This provides an interlock during the data transfer width negotiation.

The drive responds to each Initiator requested transfer width exponent as shown in the following table.

| Initiator Request | Target Response | Target Data Transfer Width |
|---|---|---|
| Ei = 00h | Et = 00h | 8 Bit Data Transfers |
| Ei > 00h | Et = 01h | 16 Bit Data Transfers |

Figure 179. Initiator Request/Target Response

If after the Target's response above the Initiator asserts the ATN signal and the first message received is either a Message Parity Error or a Message Reject message, the Target negates the data transfer width

agreement and goes to 8 bits mode. For the Message Parity Error case the implied data transfer width agreement is reinstated if the Target successfully retransmits the Wide Data Transfer Request message to the Initiator. For any other message the Target completes negotiation and goes to the negotiated data transfer width.

### 9.1.3.2 Transfer Width Negotiation started by the Target

If the Target recognizes that negotiation is required, the Target sends a Wide Data Transfer Request message to the Initiator with the transfer width exponent equal to 1 (E = 01h). The Initiator must respond by asserting the ATN signal prior to its release of ACK for the REQ/ACK handshake of the last byte of the Wide Data Transfer Request message. This provides an interlock during the wide data transfer negotiation. If the Initiator does not assert the ATN signal, the Target goes to 8 bit mode. If the Initiator does assert the ATN signal, the Target changes to the Message Out phase and receives a message from the Initiator.

If the first message received is a Wide Data Transfer Request message, the Target establishes the new data transfer mode. The drive interprets the Initiator corresponding transfer width exponent as shown in the following table.

| Initiator Request | Target Data Transfer Width |
|---|---|
| Ei = 00h | 8 Bit Data Transfers |
| Ei = 01h | 16 Bit Data Transfers |
| Ei > 01h | Send Message Reject (8 Bit Data Transfer) |

Figure 180. Target Request to Initiator

**Note:** If the corresponding transfer width exponent received from the Initiator indicates a data transfer width that is greater than 16 bits (E > 01h), the Target sends a Message Reject message to the initiator to indicate 8 bit data transfer mode.

If the first message received from the Initiator is either a Message Parity Error or a Message Reject message, the Target goes to 8 bit data transfer mode. In the case of a Message Parity Error, the wide data transfer negotiation is restarted if the Target successfully retransmits the Wide Data Transfer Request message to the Initiator.

If the first message received from the Initiator is any other message, the Target goes to 8 bit data transfer mode. The Target assumes that the Initiator does not support wide data transfer and does not attempt to renegotiate with this Initiator.

The implied agreement for wide data transfer operation is not considered to exist by the Target until the Target leaves the Message Out phase, implying that no parity error was detected. If the Target detects a parity error while attempting to receive the message from the Initiator, the Target goes to 8 bit data transfer mode. The Target will attempt to resume the wide data transfer negotiation by retrying the Message Out phase.

**Note:** If during the Message In phase of negotiations, either Target or Initiator started, ATN is asserted prior to transmission of the last byte of the message and the message is not Message Parity or Message Reject, the Target goes to 8 bit data transfer mode.

## 9.1.4 PARALLEL PROTOCOL REQUEST (01,06,04H)

Parallel Protocol Request messages are used to negotiate a synchronous data transfer agreement and a wide data transfer agreement and to set the protocol options between two SCSI devices.

| Byte | Value | Description |
|---|---|---|
| 0 | 01h | Extended message |
| 1 | 06h | Extended message length |
| 2 | 04h | Parallel Protocol Request |
| 3 | 09h | Transfer period factor (80M/T) |
| 4 | 00h | Reserved |
| 5 | X | REQ/ACK Offset |
| 6 | m | Transfer Width Exponent |
| 7 | 02h | Protocol Option (DT_REQ Only) |

Figure 181. Parallel Protocol Request.

**REQ/ACK Offset**

**X** The REQ/ACK Offset value is chosen to prevent overflow conditions in the reception buffer of the device and offset counter and the Target supports maximum offset of 63 (3Fh). A REQ/ACK Offset value of zero indicates asynchronous data transfer mode and that the Period Factor field and the Protocol Option field will be ignored.

**M** The Transfer Wide Exponent field defines the transfer width to be used during DATA IN phases and DATA OUT phases. The transfer width that is established applies to all logical units on both SCSI devices. Valid transfer widths are 8 bits (m=00h) and 16 bits (m=01h) if all the protocol options bits are zero. The only valid transfer width is 16 bits (m=01h) if any of the protocol options bits are one.

The protocol options bits (IU_REQ, DT_REQ, and QAS_REQ) are used by the originating SCSI device to indicate the protocol options to be enabled.

| **Bit** | | | | | | | |
|---|---|---|---|---|---|---|---|
| 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
| Reserved | | | | | QAS_REQ | DT_REQ | IU_REQ |

Figure 182. Bit position table for Byte 7 of Parallel Protocol Request

**QAS_REQ** (Quick Arbitrate and Select) is not supported. The bit should be zero.

**DT_REQ** A bit of zero for DT_REQ (Double Transition Enable Request) indicates that DT DATA phases are to be disabled when received from the originating SCSI device and that DT DATA phases are not supported when received from the responding SCSI device.

A DT_REQ bit of one indicates that DT DATA phases are to be enabled when received from the originating SCSI device and that DT DATA phases are supported when received from the responding SCSI device.

**IU_REQ** (Information Unit Request) is not supported. The bit should be 0.

## 9.1.5 SAVE DATA POINTER (02)

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

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

## 9.1.6 RESTORE POINTERS (03)

This message is sent from the drive to direct an initiator to copy the most recently saved pointers to the corresponding command, data, and status pointers. Command and status pointers should be restored to the beginning of the present command and status areas. The data pointer should be restored to the value at the beginning of the data area in the absence of a SAVE DATA POINTER message or to the value at the point at which the last SAVE DATA POINTER message occurred. Refer to Section 9.4, "**SCSI Bus Related Error Handling Protocol** on page 192.

## 9.1.7 DISCONNECT (04)

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

The DISCONNECT message may also be sent from the initiator to the drive to disconnect from the SCSI bus. The drive does not support the DISCONNECT message from the initiator. And it always responds by sending MESSAGE REJECT message to the initiator.

## 9.1.8 INITIATOR DETECTED ERROR (05)

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

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

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

After receiving this message the drive may retry the MESSAGE IN phase by switching to the MESSAGE IN phase with asserting REQ and repeating the previous MESSAGE IN phase.

## 9.1.9 ABORT TASK SET (06)

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

sent to the initiator for this operation. It is not an error to send the ABORT message to a logical unit that is not currently performing an operation for the initiator.

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

## 9.1.10 MESSAGE REJECT (07)

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

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

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

## 9.1.11 NO OPERATION (08)

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

## 9.1.12 MESSAGE PARITY ERROR (09)

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

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

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

## 9.1.13 LINKED TASK COMPLETE (0A)

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

## 9.1.14 LINKED TASK COMPLETE WITH FLAG (0B)

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

## 9.1.15 TARGET  RESET (0C)

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

## 9.1.16 ABORT TASK (0D)

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

## 9.1.17 CLEAR TASK SET (0E)

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

## 9.1.18 QUEUE TAG MESSAGES (20h, 21h, 22h)

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

Figure 183. Queue Tag Messages

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

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

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

### 9.1.18.1 SIMPLE (20h)

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

### 9.1.18.2 HEAD OF QUEUE (21h)

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

### 9.1.18.3 ORDERED (22h)

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

## 9.1.19 IGNORE WIDE RESIDUE

| Byte | Value | Description |
|---|---|---|
| 0 | 23H | Ignore Wide Residue message |
| 1 | 01H | Ignore |

Figure 184. Ignore Wide Residue Message Format

The Ignore Wide Residue Message is sent from the 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 (always = 01h) indicates that one byte (data bits 8-15) should be ignored. This message is sent immediately after the DATA IN phase and prior to any other messages. Even though a byte is invalid, its corresponding parity bit is valid for the value transferred.

## 9.1.20 IDENTIFY (80 - FF)

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

The IDENTIFY message is defined as follows:

**Bit 7** This bit is always set to one to distinguish the IDENTIFY message from other messages.

**Bit 6** This bit is only set to one by the initiator to grant the drive the privilege of disconnecting. If this bit is zero, the drive will not disconnect unless the initiator instructs the drive to disconnect by sending a DISCONNECT Message to the drive. This bit is set to zero when the drive sends an IDENTIFY message to the initiator.

**Bits 5-0** These bits specify the logical unit number (LUN).

Only one LUN may be identified for any one selection sequence. If the drive receives an IDENTIFY message with a new LUN after the LUN had previously been identified, the drive will go to the BUS FREE phase to signal a catastrophic error. The initiator may send more than one Identify message during a selection sequence in order to toggle disconnect/ reconnect permission if the specified LUN remains the same.

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

## 9.2 Supported Message Functions

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

- Retry SCSI Command or STATUS phase

  The retry will be caused by the following error condition.

  - The drive detected SCSI bus parity error (Command phase)
  - The drive receives INITIATOR DETECTED ERROR MESSAGE during or at the conclusion of an information transfer phase (Command Data Out or Status Phase)

  **Note:** The initiator may send the INITIATOR DETECTED ERROR message as a result of an initiator detected SCSI Bus parity error or an internal error.

- Retry MESSAGE IN phase

  The retry will be caused by the receipt of a MESSAGE PARITY ERROR message immediately following a MESSAGE IN phase.

  **Note:** The Initiator may send the MESSAGE PARITY ERROR message as a result of an Initiator detected SCSI Bus parity error during the Message In phase.

- Receipt of multiple Identify message

  - The initiator is allowed to send multiple IDENTIFY messages out in order to toggle the disconnect/reconnect permission bit. This may be used to selectively enable or disable disconnect/reconnect permission during portions of a command. Note that this function does not affect the operation of the Forced Disconnect function.

- MESSAGE REJECT during Target Disconnection

  - If the Initiator rejects the SAVE DATA POINTER message, the drive will disable disconnect/reconnect permission. This is equivalent to receiving an IDENTIFY message with bit 6 equal to zero. This will cause the drive to inhibit the pending disconnection.
  - If the initiator rejects the DISCONNECT message, the drive will not disconnect but may attempt to disconnect at a later time. This function may be used to selectively disable disconnection during portions of a command.

# 9.3 Attention Condition

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

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

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

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

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

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

# 9.4 SCSI Bus Related Error Handling Protocol

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

## 9.4.1 Unexpected BUS FREE Phase Error Condition

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

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

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

## 9.4.2 MESSAGE OUT Phase Parity Error

Depending on the model the drive will optionally retry the message phase and, if it still fails, abort the current command with CHECK CONDITION status and sense data of ABORTED COMMAND / SCSI PARITY ERROR.

## 9.4.3 MESSAGE IN Phase Parity Error (Message Parity Error)

Depending on the model the drive may retry the message phase and, if it still fails, abort the current command and go to bus free, setting sense data of ABORTED COMMAND / SCSI PARITY ERROR.

## 9.4.4 COMMAND Phase Parity Error

Depending on the model the drive may retry the command phase after sending a restore pointers message. If it still fails, it will abort the current command with CHECK CONDITION status and sense data of ABORTED COMMAND / SCSI PARITY ERROR.

## 9.4.5 DATA OUT Phase Parity Error

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

## 9.4.6 INITIATOR DETECTED ERROR Message

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

The drive will, depending on the model, optionally retry the previous phase if it is command or status. If this fails or the previous phase was a data transfer the drive will generate a CHECK CONDITION status and a Sense key of ABORTED COMMAND with additional sense code of INITIATOR DETECTED ERROR.

## 9.4.7 MESSAGE REJECT Message

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

| | |
|---|---|
| **DISCONNECT** | The drive is not disconnect but remains connected. |
| **TASK COMPLETE** | No error, continue to bus free. |
| **IDENTIFY** | Command aborted - bus freed - Sense data set to MESSAGE REJECT ERROR. |
| **LINKED CMD CMPLT** | Command aborted - link broken - bus freed - sense data set to MESSAGE REJECT ERROR. |
| **MESSAGE REJECT** | Command aborted - STATUS phase executed with CHECK CONDITION - sense data set to MESSAGE REJECT ERROR. |
| **RESTORE POINTERS** | Command aborted - status set to CHECK CONDITION - sense will be set with the error that caused the RESTORE POINTERS message to be issued. (Assuming that error recovery is in progress) |
| **SAVE DATA POINTER** | The drive will not disconnect from the SCSI bus and tt will not be considered an error. |
| **No previous Msg** | The command is aborted, the bus freed, and Sense data is set to MESSAGE REJECT ERROR. This occurs when the drive has not sent a message, but gets a MESSAGE REJECT from the initiator. |

This page intentionally left blank.

# 10.0 Additional information

This chapter provides additional information or descriptions of various functions, features, or operating models supported by the Target that are not fully described in previous chapters.

## 10.1 SCSI Protocol

There are various operating conditions that prevent the Target from executing a SCSI command. This section describes each of these operating conditions and their relative priority.

### 10.1.1 Priority of SCSI Status Byte Reporting

After establishing the I_T_L nexus or I_T_L_Q nexus the Target must first determine whether command execution is allowed. Execution is deferred until a later time if the command must be added to the command queue. Execution may also be prevented by an internal Target condition that requires the reporting of a Check Condition, Queue Full, Busy, or Reservation Conflict Status. There are several different internal conditions to be active at the same time. The order in which the Target checks for each of these conditions determines their priority (highest priority first) as follows:

1. Check Condition status for invalid Logical Unit Number. (See Section 10.1.2 "Invalid LUN in Identify Message" 195)
2. Check Condition status for Incorrect Initiator Connection (See Section 10.1.3 “Incorrect Initiator Connection” on page 196)
3. Busy Status or Queue Full Status (See 10.1.4 “Command Processing During Execution of Active I/O process” on page 196)
4. Check Condition status for Unit Attention condition (See Section 10.1.5 “Unit Attention Condition” on page 198)
5. Check Condition status during Startup and Format operations (See Section 10.1.6 “Command Processing During Startup and Format Operations” on page 199)
6. Check Condition status for Deferred Error Condition (See Section 10.1.8 “Deferred Error” on page 199)
7. Reservation Conflict status (See Section 10.1.10 “Command Processing While Reserved” on page 206)
8. Check Condition status for invalid command opcode
9. Check Condition status for invalid command descriptor block

The highest priority internal condition that prevents command execution is reported by the Target provided there is no bus error.

For all Check Conditions Sense data is built by the target provided a valid LUN address is known. Sense data is cleared by the Target upon receipt of any subsequent command to the LUN from the initiator receiving the Check Condition.

### 10.1.2 Invalid LUN in Identify Message

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

### 10.1.2 .1 Case 1 - Selection message sequence with Inquiry command

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

### 10.1.2 .2 Case 2 - Selection message sequence with any other command

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

### 10.1.2 .3 Case 3 - After selection message sequence

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

## 10.1.3 Incorrect Initiator Connection

Incorrect Initiator Connection error is reported if any of the following conditions occur:

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

  **Note:** It is not an Incorrect Initiator Connection to send a command without a Queue tag message when sense is pending on the logical unit for the Initiator that issues the Request Sense command. (If the command is not Request Sense or Inquiry, sense data is cleared upon receipt of the command. )
- an Initiator attempts to establish an I_T_L_Q nexus when an I/O process (either queued or active) with an I_T_L_Q nexus already exists from a previous connection with the same initiator.

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

## 10.1.4 Command Processing during execution of active I/O process

When the Target is not executing any active I/O processes, a new I/O process is permitted to execute (unless execution is prevented by another internal Target condition listed in Section 10.1.1 “Priority of SCSI Status Byte Reporting” on page 195).

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

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

- Queue Full Status is returned
- Busy Status is returned

If an active I/O process does exist when the Target receives a new command, the Target determines how the new command should be handled based on the following rules:

- Check Condition Status is returned with Sense Key set to Aborted Command for an Overlapped Commands Attempted error if
  - See 10.1.3 “Incorrect Initiator Connection” on page 196.
- the command is permitted to execute if the command is an Inquiry or Request Sense command
- Check Condition Status is returned with Sense Key set to Logical Unit Not Ready if the startup operation or format operation is an active process.
- the command is permitted to execute if the conditions to execute concurrently are met. (See Section 10.5 “Concurrent I/O Process.” on page 209 )
- the command is added to the command queue for an I_T_L nexus if all the following conditions exist:
  - no Queue Tag message was received during the connection which established the I/O process
  - disconnection is allowed for the current I/O process
  - there is no queued I/O process or active I/O process corresponding to the I_T_L nexus for the current I/O process
  - the command is not linked to a previous command
- the command is added to the command queue for an I_T_L_Q nexus if the following conditions exist:
  - a Queue Tag message was received during the connection which established the I/O process
  - Tagged Queuing is enabled (DQue = 0)
  - an I/O process (either active or queued) exists at the Target for this Initiator
  - disconnection is allowed for the current I/O process
  - there is no queued I/O process or active I/O process corresponding to the I_T_L_Q nexus for the current I/O process
  - the command is not linked to a previous command.
- Queue Full Status is returned if any one of the following conditions exists:
  - the command would otherwise be queued (according to the rules described above) but the command queue is full and all slots are utilized
  - the command would otherwise be queued (according to the rules described above) but all of the available command queue slots not reserved for use by another initiator are utilized
  - Tagged Queuing is enabled (DQue = 0) and a Format Unit command was previously queued but has not yet begun execution
  - Tagged Queuing is enabled (DQue = 0) and a Unit start command was previously queued but has not yet begun execution
- Busy Status is returned if any of the following conditions exists:
  - Tagged Queuing is disabled (DQue = 1) and a Format Unit command was previously queued but has not yet begun execution
  - Tagged Queuing is disabled (DQue = 1) and a Unit start command was previously queued but has not yet begun execution
  - the command would otherwise be queued (according to the rules described above) but disconnection is not allowed for the current I/O process.

If a command is queued, command execution may still be prevented at a later time when the command is dequeued to become an active I/O process. This occurs if command execution is prevented at the time the command is dequeued by another internal Target condition listed in Section 10.1.1 "Priority of SCSI Status Byte Reporting" on page 195.

## 10.1.5 Unit Attention Condition

The drive will generate a unit attention condition for each initiator whenever these conditions are in effect:

- The drive has been reset. This includes Power On Reset, SCSI Bus Reset, SCSI TARGET RESET message.
- The Mode parameters in effect for this initiator have been changed by another initiator.
- A Log Select was issued by an another initiator thus clearing log parameters.
- The microcode has been changed. WRITE BUFFER command has been executed to download microcode. In this case a unit attention condition is generated for all initiators except the one that issued the command.
- Commands are cleared by another initiator. This condition is generated against the initiator that has queued commands if
  - a Clear Queue Message is received.
  - a Contingent Allegiance Condition is cleared when QERR (in Mode Page 0A) is 1.
  - a DQue is set to 1 while queued command exist.
- SCSI Bus tranceiver mode change (LVD <==> Single Ended).
- Reservations or Registrations were preempted by another initiator.
- A Predictive Failure Analysis thershold has been reached.

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

If the drive receives a command from each initiator before reporting a CHECK CONDITION status for a pending unit attention condition for that initiator, the response of the drive varies with the command as follows:

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

**REQUEST SENSE** If the drive has an available pending sense data for the initiator, the drive sends the pending sense data and preserves the unit attention condition for the initiator. If the drive does not have an available pending sense data for the initiator, the drive sends sense data for the unit attention condition and clears the unit attention condition for the initiator.

**ALL OTHER** The drive terminates the command with a CHECK CONDITION status and preserves the unit attention condition.

If the drive receives a command from each initiator after reporting a CHECK CONDITION status for a pending unit attention condition for that initiator, the response varies with the command as follows:

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

**ALL OTHER** The drive executes the command with GOOD status and clears the unit attention condition unless another unit attention condition exists. And then the sense data for the unit attention condition is lost.

## 10.1.6 Command Processing During Startup and Format Operations

If the Target receives a command from an Initiator while the Target is executing a startup or format operation, the response of the Target varies with the command as follows:

**INQUIRY** The drive sends inquiry data and returns appropriate status.

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

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

For the START/UNIT STOP and the Auto-start operation, the qualifier returned is LOGICAL UNIT IS IN PROCESS OF BECOMING READY. For the FORMAT UNIT command, the qualifier returned is LOGICAL UNIT NOT READY, FORMAT IN PROGRESS, and the Sense key specific bytes are set to return the progress indication.

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

## 10.1.7 Internal Error Condition

The Target generates an Internal Error condition for all Initiators when an internally initiated operation ends with an unrecoverable error, that is, the startup sequence for Auto Start enabled terminates after the SCSI bus has been enabled and prior to completion of the bring-up sequence.

An Internal Error condition causes Sense data to be generated and saved for all Initiators. The Error Code field of the Sense is set for a Current Error (70h) and the Sense Key is set to HARDWARE ERROR. Recovered errors are not reported.

The Internal Error condition persists for each Initiator until that Initiator clears the condition from the logical unit as described below. Several commands are handled as special cases during an Internal Error condition. These cases are also discussed.

If the Target receives a command from an Initiator while an Internal Error condition exists for that Initiator, the response of the Target varies with the command as follows:

**INQUIRY** The drive executes the command with GOOD status and does not clear the Internal Error condition.

**REQUEST SENSE** The drive executes the command, returns the sense data generated by the Internal Error condition, returns Good Status, and clears the Internal Error condition for that Initiator.

**ALL OTHER** The drive terminates the command with a CHECK CONDITION status and clears the Internal Error condition.

## 10.1.8 Deferred Error

Error code (71h) of sense data indicates that the Check Condition status returned is the result of an error or exception condition that occurred during execution of a previous command for which Good status has already been returned.

The drive creates an Deferred Error condition when

- Execution of a Format Unit command with the immediate bit of one ends with an error.
- Execution of a Write command with WCE (Write Cache Enable) bit of one ends with an error.

## 10.1.9 Degraded Mode

There are certain errors or conditions which may impair the ability of the drive to function normally. Rather than fail hard the drive is designed to be as responsive as possible. Also, in most cases, some action on the part of the initiator may be used to restore normal operation. This mode of limited operation is called Degraded Mode.

There are 3 conditions in the Degraded Mode:

- Spindle Motor Degrade which is caused by one of the following conditions:
  - Spindle Motor is not started by the option jumper setting (Disable Auto Spin Up)
  - Spindle Motor is delayed from spinning up by the option jumper setting (Auto Start Delay)
  - Spindle Motor was started (by POR or Unit Start command) and the Target is under Self Configuration.
  - Spindle Motor Failed to start.
  - Spindle Motor was stopped by Unit Stop command after the Target successfully completed the Self Configuration.

- Self Configuration Failure Degraded which is caused by one of the following conditions:
  - RAM Code, Configuration Sector Read Failure
  - RAM Code, Configuration Sector Revision Mismatch

- Format Command Failure Degraded. This condition is caused when Format Unit command failed or was interrupted abnormally (Mode Page 0, byte 5, bit 4 FDD controls Format Degraded mode)

### 10.1.9 .1 Response to SCSI Command in Degraded Mode - Disable Auto Start

The tables on the following pages show the degraded mode status with acceptable commands and additional sense codes

| **Command (w/Option)** | **Response** |
|---|---|
| Request Sense | Executed. The Target may return Sense Key 02h (Not Ready) ASC/ASCQ 0402h (Initialize Command Required) |
| Inquiry (EVPD=0) | Executed |
| Inquiry (EVPD=1) | Executed and Check Condition is returned with Sense Key 05h (Illegal Request) ASC/ASCQ 2400h (Invalid Field in CDB) |
| Test Unit Ready | Executed and Check Condition is returned with Sense Key 02h (Not Ready) ASC/ASCQ 0402h (Initialize Command Required) |
| Start Stop Unit (Start) | Executed<br>- Success: Good Status is returned. Motor Degraded Mode is cleared<br>- Spindle Motor Start Failure: Check Condition with Sense Key 02h (Not Ready) ASC/ASCQ 0400h (Start Spindle Motor Fail)<br>- Self Configuration Failure: Check Condition with Sense Key 02h (Not Ready) ASC/ASCQ 4080h (Diag Fail-Bring up Fail)<br>Sense Key 02h (Not Ready) ASC/ASCQ 4085h (Diag Fail-RAM Code NOT load) |
| Start Stop Unit (Stop) | Executed. Good Status is returned. Motor Degraded Mode is NOT cleared |
| Other Commands | Not Executed. Check Condition Status is returned with Sense Key 02h (Not Ready) ASC/ASCQ 0402h (Initialize Command Required) |

Figure 185. Spindle Motor Degraded Mode - Disable Auto Start

### 10.1.9 .2 Response to SCSI Command in Degraded Mode - Auto Start Delay/ Spinning Up

| **Command (w/Option)** | **Response** |
|---|---|
| Request Sense | Executed. The Target may return Sense Key 02h (Not Ready) ASC/ASCQ 0401h (In Process of Becoming Ready) |
| Inquiry (EVPD=0) | Executed |
| Inquiry (EVPD=1) | Executed and Check Condition is returned with Sense Key 05h (Illegal Request) ASC/ASCQ 2400h (Invalid Field in CDB) |
| Test Unit Ready | Executed and Check Condition is returned with Sense Key 02h (Not Ready) ASC/ASCQ 0401h (In Process of Becoming Ready) |
| Start Stop Unit (Start) | Executed<br>- Success: Good Status is returned. Motor Degraded Mode is cleared<br>- Spindle Motor Start Failure: Check Condition with Sense Key 02h (Not Ready) ASC/ASCQ 0400h (Start Spindle Motor Fail)<br>- Self Configuration Failure: Check Condition with Sense Key 02h (Not Ready) ASC/ASCQ 4080h (Diag Fail-Bring up Fail)<br>Sense Key 02h (Not Ready) ASC/ASCQ 4085h (Diag Fail-RAM Code NOT load) |
| Other Commands | Not Executed. Check Condition Status is returned with Sense Key 02h (Not Ready) ASC/ASCQ 0401h (In Process of Becoming Ready) |

Figure 186. Spindle Motor Degraded Mode - Auto Start Delay/Spinning Up

### 10.1.9 .3 Response to SCSI Command in Degraded Mode - Spindle Start Failure

| **Command (w/Option)** | **Response** |
|---|---|
| Request Sense | Executed. The Target may return Sense Key 02h (Not Ready) ASC/ASCQ 0400h (Start Spindle Motor Fail) |
| Inquiry (EVPD=0) | Executed |
| Inquiry (EVPD=1) | Executed and Check Condition is returned with Sense Key 05h (Illegal Request) ASC/ASCQ 2400h (Invalid Field in CDB) |
| Test Unit Ready | Executed and Check Condition is returned with Sense Key 02h (Not Ready) ASC/ASCQ 0400h (Start Spindle Motor Fail) |
| Start Stop Unit (Start) | Executed<br>- Success: Good Status is returned. Motor Degraded Mode is cleared<br>- Spindle Motor Start Failure: Check Condition with Sense Key 02h (Not Ready) ASC/ASCQ 0400h (Start Spindle Motor Fail)<br>- Self Configuration Failure: Check Condition with Sense Key 02h (Not Ready) ASC/ASCQ 4080h (Diag Fail-Bring up Fail)<br>Sense Key 02h (Not Ready) ASC/ASCQ 4085h (Diag Fail-RAM Code NOT load) |
| Start Stop Unit (Stop) | Executed. Good Status is returned. Motor Degraded Mode is NOT cleared |
| Other Commands | Not Executed. Check Condition Status is returned with Sense Key 02h (Not Ready) ASC/ASCQ 0400h (Start Spindle Motor Fail) |

Figure 187. Spindle Motor Degraded Mode - Spindle Start Failure

### 10.1.9 .4 Response to SCSI Command in Degraded Mode - Spindle Stopped by Unit Stop Command

| **Command (w/Option)** | **Response** |
|---|---|
| Request Sense | Executed. The Target may return Sense Key 02h (Not Ready) ASC/ASCQ 0402h (Initialize Command Required) |
| Inquiry (EVPD=0) | Executed |
| Inquiry (EVPD=1) | Executed |
| Test Unit Ready | Executed and Check Condition is returned with Sense Key 02h (Not Ready) ASC/ASCQ 0402h (Initialize Command Required) |
| Start Stop Unit (Start) | Executed<br>- Success: Good Status is returned. Motor Degraded Mode is cleared<br>- Spindle Motor Start Failure: Check Condition with Sense Key 02h (Not Ready) ASC/ASCQ 0400h (Start Spindle Motor Fail)<br>- Self Configuration Failure: Check Condition with Sense Key 02h (Not Ready) ASC/ASCQ 4080h (Diag Fail-Bring up Fail)<br>Sense Key 02h (Not Ready) ASC/ASCQ 4085h (Diag Fail-RAM code NOT load) |
| Start Stop Unit (Stop) | Executed. Good Status is returned. Motor Degraded Mode is NOT cleared |
| Other Commands | Not Executed. Check Condition Status is returned with Sense Key 02h (Not Ready) ASC/ASCQ 0402h (Initialize Command Required) |

Figure 188. Spindle Motor Degraded Mode - Spindle Stopped by Unit Stop Command

### 10.1.9 .5 Self Configuration Failure Degraded Mode

| **Command (w/Option)** | **Response** |
|---|---|
| Request Sense | Executed. The Target may return<br>Sense Key 02h (Not Ready) ASC/ASCQ 4080h (Diag Fail-Bring up Fail)<br>Sense Key 02h (Not Ready) ASC/ASCQ 4085h (Diag Fail-RAM code NOT load) |
| Inquiry (EVPD=0) | Executed |
| Inquiry (EVPD=1) | Executed and Check Condition is returned with<br>Sense Key 05h (Illegal Request) ASC/ASCQ 2400h (Invalid Field in CDB) |
| Test Unit Ready | Executed and Check Condition is returned with<br>Sense Key 02h (Not Ready) ASC/ASCQ 4080h (Diag Fail-Bring up Fail)<br>Sense Key 02h (Not Ready) ASC/ASCQ 4085h (Diag Fail-RAM code NOT load) |
| Start Stop Unit (Start) | Executed<br>- Success: Good Status is returned. Motor Degraded Mode is cleared<br>- Spindle Motor Start Failure: Check Condition with<br>Sense Key 02h (Not Ready) ASC/ASCQ 0400h (Start Spindle Motor Fail)<br>- Self Configuration Failure: Check Condition with<br>Sense Key 02h (Not Ready) ASC/ASCQ 4080h (Diag Fail-Bring up Fail)<br>Sense Key 02h (Not Ready) ASC/ASCQ 4085h (Diag Fail-RAM code NOT load) |
| Write Buffer (Download and Save) | Executed.<br>- Success: Good Status is returned. Motor Degraded Mode is cleared<br>- Self Configuration Failure: Check Condition with<br>Sense Key 02h (Not Ready) ASC/ASCQ 4080h (Diag Fail-Bring up Fail)<br>Sense Key 02h (Not Ready) ASC/ASCQ 4085h (Diag Fail-RAM code NOT load) |
| Other Commands | Not Executed. Check Condition Status is returned with<br>Sense Key 02h (Not Ready) ASC/ASCQ 4080h (Diag Fail-Bring up Fail)<br>Sense Key 02h (Not Ready) ASC/ASCQ 4085h (Diag Fail-RAM code NOT load) |

Figure 189. Self Configuration Failure Degraded Mode

### 10.1.9 .6 Format Command Failure Degraded Mode

| Command (w/Option) | Response |
|---|---|
| Request Sense | Executed. The Target may return<br>Sense Key 02h (Not Ready) ASC/ASCQ 3100h (Format Corrupted)<br>Sense Key 03h (Medium Error) ASC/ASCQ 3100h (Format Corrupted) |
| Inquiry (EVPD=0) | Executed |
| Inquiry (EVPD=1) | Executed |
| Test Unit Ready | Executed and Check Condition is returned with Sense Key 02h (Not Ready) ASC/ASCQ 3100h (Format Corrupted) |
| Format Unit | Executed<br>- Success: Good Status is returned. Format Degraded Mode is cleared<br>– Failure: Check Condition Status is returned and Format Degraded Mode is NOT cleared. |
| Other Commands | Not Executed. Check Condition Status is returned with Sense Key 03h (Medium Error) ASC/ASCQ 3100h (Format Corrupted) |

Figure 190. Format Command Failure Degraded Mode

Note: Mode Page 0 byte 5 bit 4 (FDD) = 0

## 10.1.10 Command Processing while Reserved

A logical unit is reserved after successful execution of the Reserve command. Each time a Reserve command is executed successfully, the Target records the SCSI ID of the Initiator that made the reservation and the SCSI ID of the Initiator that is to receive the reservation. This information is needed to determine whether subsequent commands should be permitted or if the Reservation Conflict Status should be reported. The Initiator that made the reservation is the Initiator that issued the Reserve command. The Initiator to receive the reservation may be either the same or a different Initiator (third-party reservation).

If the logical unit is reserved when a new command is received, the Target examines the command opcode and the SCSI ID of the issuing Initiator to determine whether a Reservation Conflict Status should be returned based on the following rules:

- If the issuing Initiator is the one that made the reservation and also the one to receive the reservation, then all commands are permitted.
- If the issuing Initiator is neither the one that made the reservation nor the one to receive the reservation, then
  - A Request Sense or Inquiry command is permitted.
  - A Release command is permitted but is ignored.
  - Any other command results in a Reservation Conflict Status.
- If the issuing Initiator is the one that made the reservation but is not the one to receive the reservation, then
  - An Inquiry, Request Sense, Reserve, or Release command is permitted.

  - Any other command results in a Reservation Conflict Status.
- If the issuing Initiator is not the one that made the reservation but is the one to receive the reservation, then
  - A Reserve command results in a Reservation Conflict Status.
  - A Release command is permitted but is ignored.
  - Any other command is permitted.

If a Reservation Conflict Status is not reported and the command is permitted, then the Target checks the next highest priority internal condition to determine whether execution is allowed. See Section 10.1.1 “Priority of SCSI Status Byte Reporting” on page 195.

# 10.2 Priority Commands

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

- Inquiry
- Request Sense

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

These commands are never queued whether or not the command is sent with a queue tag. However, the rule for an Incorrect Initiator Connection still apply to priority commands. See 10.1.3 “Incorrect Initiator Connection” on page 196.

# 10.3 Command Queuing

When the initiator specifies that the drive shall disable command queuing, the initiator must send only untagged commands. When the initiator specifies that the target shall enable command queuing, the initiator may send either tagged or untagged command, but shall not use both at the same time.

The following commands are never queued and will be immediately executed without Bus disconnection:

- Priority Commands (i.e.: Request Sense and Inquiry)
- Commands linked to previous commands. These are defined to be part of a single I/O process. (Linked commands are always executed immediately following the previous command from the same initiator. No other Initiator's command are allowed to be executed between two linked commands.)
- Commands for which disconnection is not allowed. (These may result in a Busy Status.)
- Commands in which a SCSI bus error occurred between selection and first disconnection following the receipt of the CDB.
- Commands for an invalid LUN.
- Commands which cause an OVERLAPPED COMMANDS ATTEMPTED error (see Section 11.1.3, “Incorrect Initiator Connection”).

## 10.3.1 Queue depth

Any initiator can queue at least one command at any time irrespective of the actions of any other initiators in the system. A single initiator may queue up to 128 commands, if no other initiator has more than

one command in the queue, although at times this maximum may be reduced as the drive can reserve command blocks for internal use.

## 10.3.2 Tagged queuing

Commands with a tag message are saved in the command queue. Queued commands will be reordered by the target defined rule. See the Section 10.4, “Command reordering” on page 209 for details.

## 10.3.3 Untagged queuing

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

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

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

## 10.3.4 Command queuing rule

Commands can be received during an active I/O process if the Bus is free. If CPE (concurrent process enable) bit of Mode page 0 is 0, only a single command except 2 commands (Inquiry and Request Sense) can be executed at the same time. In that case, a command starts to be executed after the previous command has completed. If CPE (concurrent process enable) bit of Mode page 0 is 1, I/O processes of Read(6), Read extend(10), Write(6) and Write extend(10) can be active at the same time. See Section 10.5 **Concurrent I/O Process**" on page 209 for details.

## 10.3.5 Queue Full status

This status is returned when a Simple, ORDERED or HEAD OF QUEUE message is received and the command queue is full. The I/O process is not placed in the command queue. Since one queue element is reserved for each initiator, any untagged command that does not cause Incorrect Initiator Connection will not cause Queue Full status.

## 10.3.6 Device behavior on Command queuing

1. Initiators must send a Queue tag immediately after the Identify message in Message Out phase just after Selection. Targets send a Simple message immediately after the Identify message in Message In phase just after Reselection.
2. Each initiator can issue either a tagged command or an untagged command exclusively at the same time. Other initiators can exist which operate mutually exclusively with tagged or untagged commands
3. When DQue (Disable queue) of mode page 0Ah is 1, if an initiator issues a tagged command, the drive returns "Message Reject" message (07h) and receives that command as an untagged command.
4. Queue Tag number does not affect the order of execution.
5. If an initiator issues a command with a queue tag which is the same as the current I/O process or queued I/O process. The target returns Incorrect Initiator connection.
6. A series of linked commands are a single I/O process and are assigned the queue tag established in the initial selection. A command received with a Head-of-Queue tag message shall not suspend a series of linked commands for which the target has begun execution.

7. If DQue is changed to 1 while queued commands exist, all queued commands for the all initiators will be aborted. All future commands received from any initiator with a queue tag will be processed as untagged commands with a message reject message being returned immediately after the qtag is received by the target.

## 10.4 Command reordering

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

## 10.5 Concurrent I/O Process

The Concurrent I/O process when multiple I/O processes are active (not queued) on the same logical unit at the same time. The target may start the data phase of an I/O process while another I/O process is not completed. The following I/O processes are allowed to execute concurrently:

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

When an I/O process ends in Check Condition Status, the drive enters the Contingent Allegiance Condition and other queued I/O processes from all initiators on the same logical unit will not reconnect and will not complete the execution until the sense data is cleared. See Section 11.16, “Contingent allegiance Condition,” for details. If an I/O process (P-1) encounters an error while another I/O process (P-2) is active, the drive returns Check Condition to P-1 and P-2. The drive may continue P-2 until its convenient point to suspend, but may not send a Status. After the initiator clears the Contingent Allegiance condition, the drive will resume or terminate P-2 according to QErr bit of Control mode page.

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

## 10.6 Back to Back Write

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

## 10.7 Write Cache

If the WCE (Write cache enable) bit is 1, the drive returns Good Status and Task complete message and goes to Bus Free immediately after receiving the data of the last sector before actually writing the data onto the media.

If the drive detects an error after it returns a Good Status, the drive sets a Deferred Error (Error Code of sense data = 71h) and a following command will be returned with Check Condition and the Contingent

allegiance condition is established. Under the Contingent allegiance condition all queued processes including commands from other initiators are suspended.

## 10.8 Automatic Rewrite/Reallocate

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

Automatic and Recommend Reallocate operate from within the read/write command. When an automatic reallocation occurs, the read or write command takes longer to complete.

This operation is sometimes referred to as auto-reassignment due to its similarity to the operation performed by the reassign command.

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

**ARRE=1:** An error site determined to need rewriting or reallocation during a read is automatically rewritten or reallocated at the conclusion of the read and prior to the sending of the status. The site will be automatically rewritten or reallocated only if the data has been successfully read.

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

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

- Read(6)
- Read(10)
- Verify
- Verify Portion of Write and Verify

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

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

**AWRE=1 :** An error site determined to need reassignment during a write is automatically reallocated at the conclusion of the write and prior to sending the status. The site will be automatically reallocated only if the write recovery succeeded at the conclusion of the write.

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

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

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

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

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

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

| Key | Code | Qual | Description |
|---|---|---|---|
| 1 | 17 | 01 | Recovered Data with retries |
| 1 | 17 | 06 | Recovered Data without ECC - Auto Reallocated |
| 1 | 17 | 07 | Recovered Data without ECC - Recommend Reassign |
| 1 | 17 | 09 | Recovered Data without ECC - Data Rewritten |
| 1 | 18 | 00 | Recovered Data with ECC |
| 1 | 18 | 02 | Recovered Data with ECC - Auto Reallocated |
| 1 | 18 | 05 | Recovered Data with ECC - Recommend Reassign |
| 1 | 18 | 07 | Recovered Data with ECC - Data Rewritten |

Figure 191. Sense data combinations with auto/recommend rewrite/reallocate

# 10.9 Segmented Caching

## 10.9.1 Overview

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

## 10.9.2 Read Ahead

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

The drive initiates the Read ahead function when the following conditions exist:

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

If SCSI reset or target reset message is received, all contents of segmented buffer is flushed.

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

# 10.10 Reselection Timeout

If reselection fails, it will be retried one or more times depending on the drive model. Please see the individual drive specifications for the number of retries allowed.

# 10.11 Single Initiator Selection

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

# 10.12 Non-arbitrating systems

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

## 10.13 Selection without ATN

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

Phase retries may be allowed if a subsequent Identify message is received.

## 10.14 Multiple Initiator Environment

### 10.14.1 Initiator Sense Data

Separate sense data is reserved for each I-T-L. Each sense data is maintained independent of commands from other initiators.

### 10.14.2 Initiator Mode Select/Mode Sense Parameters

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

### 10.14.3 Initiator Data Transfer Mode Parameter

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

## 10.15 Contingent Allegiance Condition

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

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

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

# 10.16 Reset

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

This device uses the Hard reset option as defined in the SCSI-3 standard.

## 10.16.1 Reset Sources

There are four sources of resets detected by the target:

| Reset Name | Reset Source |
|---|---|
| Power-On Reset | This is the signal generated by the hardware at initial power-on |
| Self-Initiated reset | This is a software-generated reset that occurs when a catastrophic error is detected by the microcode. |
| SCSI Bus Reset | This is a reset generated when the SCSI bus control line RST goes active. |
| Target Reset Message | This is the reset generated by the SCSI Target Reset Message (0Ch). |

## 10.16.2 Reset Actions

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

### 10.16.2 .1 Power-On reset and Self-Initiated reset

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

- a power-up sequence
- a startup sequence is necessary to put the drive in a ready state

### 10.16.2 .2 SCSI Bus reset and SCSI Target Reset message

These two reset conditions cause the following to be performed.

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

  **Note:** The power-up sequence, having already completed, is not rerun.
- If reset occurs while a physical sector is being written, the write operation is disabled after the current physical sector is written. Data is not lost as long as power stays valid until the physical sector being written is completed.

# 10.17 Diagnostics

The drive will execute Power on Diagnostics at power on time to assure the correct operation of the drive by validating components (ROM, RAM, Sector Buffer, EEPROM, HDC, Spindle Motor, Actuator), checking stored information in the Reserved Area and EEPROM, and verifying fault detects circuits.

Self-test can be invoked by issuing a SEND DIAGNOSTIC command.

## 10.17.1 Power on Diagnostics

At power on time the following tests are executed:

1. Validation of ROM and EEPROM
2. RAM test for internal RAM
3. Test and Initialize HDC registers
4. RAM test for Sector Buffer
5. Start Spindle Motor (if Auto spin up enable)
6. Calibration of Actuator
7. Read/Write test for all Heads
8. Validation of RAM code and data table (RDM, Log, Mode Page) from the Reserved Area

If Auto spin up is disabled, steps 5 - 8 will be executed when the first START UNIT STOP command with Start bit.

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

Faults detected after the successful completion of the HDC test section will be reported as CHECK CONDITION status to the Initiator on the first command issued after a fault is detected (except for the INQUIRY, REPORT LUNS and REQUEST SENSE commands). The INQUIRY, REPORT LUNS and REQUEST SENSE commands will always be responded with a GOOD status. Detecting a fault during power on will not terminate execution of the tests nor will it terminate the power on process.

## 10.17.2 Self-test via SEND DIAGNOSTIC Command

### 10.17.2 .1 Default Self-test

The default self-test is invoked by the SlfTst bit in the SEND DIAGNOSTIC command. The response is simply a GOOD status if the test is successful or a CHECK CONDITION status if the test fails.

The following tests are performed by the default self-test (in the order defined):

1. **Spin check** is to check if the spindle motor is running at the correct speed.
2. **Write, Read and Compare test** is a disk read/write test. It writes data to a predefined location in the reserved area and then reads it back and validates the content. Head 0 is used for this test.
3. **ECC circuit test** is a test for ECC circuit to ensure that errors can be corrected by the circuit.
4. **Seek test** is a servo test. It validates seeks to 256 random locations out of the full volume.

### 10.17.2 .2 Short and Extended Self-tests

There are two other types of self-tests that may be invoked using the Function Code field in the SEND DIAGNOSTIC command: a short self-test and an extended self-test. The tests performed in the short and extended self-tests are described later. The time required by a logical unit to complete its extended self-test is specified in the Extended self-test Completion Time field in the Control Mode Page. The results of self-test can be retrieved via the LOG SENSE command for log page 10.

### 10.17.2 .3 Self-test Modes

There are two modes for short and extended self-tests: a foreground mode and a background mode. These modes are described in the following clauses.

**Foreground mode**
When the drive receives a SEND DIAGNOSTIC command specifying a self-test to be performed in the foreground mode, the drive will return status for that command after the self-test has been completed. While performing a self-test in the foreground mode, the drive will respond to all commands except INQUIRY, REPORT LUNS, and REQUEST SENSE with a CHECK CONDITION status, a sense key of NOT READY and an additional sense code of LOGICAL UNIT NOT READY - SELF-TEST IN PROGRESS.

If the drive is performing a self-test in the foreground mode and a test error occurs, the drive will update the self-test results log page and report CHECK CONDITION status with a sense key of HARDWARE ERROR and an additional sense code of LOGICAL UNIT FAILED SELF-TEST. The application client may obtain additional information about the failure by reading the self-test results log page.

An application client may terminate a self-test that is being performed in the foreground mode using an ABORT TASK, ABORT TASK SET, or CLEAR TASK SET task management function. If the drive receives an ABORT TASK, ABORT TASK SET, or CLEAR TASK SET task management function while performing a self-test in the foreground mode, it will abort the self-test and update the self-test results log page.

**Background mode**
When the drive receives a SEND DIAGNOSTIC command specifying a self-test to be performed in the background mode, the drive will return status for that command as soon as the command descriptor block has been validated. After returning status for the SEND DIAGNOSTIC command specifying a self-test to be performed in the background mode, the drive will initialize the self-test results log page as follows. The Function Code from the SEND DIAGNOSTIC command will be placed in the Function Code field in the log page. The self-test Results field shall be set to 0Fh. After the self-test results log page is initialized, the drive will begin the first self-test segment.

While the device server is performing a self-test in the background mode, it shall terminate with a CHECK CONDITION status any SEND DIAGNOSTIC command it receives that meets one of the following criteria:

a) The SlfTst bit is one

b) The Function Code field contains a value other than 000b or 100b.

When terminating the SEND DIAGNOSTIC command, the sense key shall be set to NOT READY and the additional sense While performing a self-test in the background mode, the drive will suspend the self-test to service any other command other than SEND DIAGNOSTIC (with Function Code field set to 100b) WRITE BUFFER (with the mode set to any download microcode option), FORMAT UNIT and START UNIT STOP command. Suspension of the self-test to service the command will occur within 2 seconds. If SEND DIAGNOSTIC (with Function Code field set to 100b), WRITE BUFFER (with the mode set to any download microcode option), FORMAT UNIT or START UNIT STOP command is received, the drive will abort the self-test, update the self-test log, and service the command within two seconds after the command descriptor block has been validated.

An application client may terminate a self-test that is being performed in the background mode by issuing a SEND DIAGNOSTIC command with the Function Code field set to 100b (Abort background self-test function).

**Elements common to foreground and background self-test modes**
The Progress Indication field returned in response to a REQUEST SENSE command may be used by the application client at any time during execution of a self-test to poll the progress of the test. While executing a self-test unless an error has occurred, the drive will respond to a REQUEST SENSE

command by returning a sense key of NOT READY and an additional sense code of LOGICAL UNIT NOT READY - SELF-TEST IN PROGRESS with the sense key specific bytes set for progress indication. The application client may obtain information about the twenty most recently completed self-tests by reading the self-test results log page. This is the only method for an application client to obtain information about self-tests performed in the background mode. The default self-test results are not logged in the log page.

**Tests performed in the Short and Extended Self-test**

The following table defines the tests performed in the short and extended self test. They are defined by their segment number which is also used to report Self-Test Results, in Log Sense Page 10. Note that the only difference between the Short and the Extended tests, is the sequential verify test in segment 9h. Also note that either of these tests can be run in foreground or background mode as previously described.

| Segment Number | Short Self-Test | Extended Self-Test | Test Description |
|---|---|---|---|
| 1h | Drive Ready Test | | Internal check to insure drive is "ready", similar to a Test Unit Ready command. |
| 2h | RAM Test | | Write/Read/Compare of drive's entire sector buffer |
| 3h | Drive Diagnostics | | This test is comprised of the Default Self Test as defined in Section 10.17.2 .1 |
| 4h | SMART | | Perform SMART testing and check results to ensure that SMART threshold criteria are not exceeded |
| 5h | Low Level Format check | | Check to insure that the media is currently not in the MEDIA FORMAT CORRUPTED state. |
| 6h | Physical Head Check | | Write/Read test on each head in a predefined location in the drive's Reserved Area of the disk. |
| 7h | Error Log pages check | | Check log pages 2, 3, and 5; report if any uncorrectable error counters are > 0 |
| 8h | Random Verify | | Perform 4000 random verify operations and insure no uncorrectable errors. |
| 9h | - Verify First 300MB<br>- Verify Last 100 MB | Verify all LBA's | Sequential verify operation. Ensure that no uncorrectable errors occur within the verify range. |
| Ah | Verify sectors adjacent to P & G list | | Perform Read commands on sectors that are adjacent to Glist entries and Plist entries and determine that there are no uncorrectable errors. |
| Bh | Recheck Error Log pages | | Same as segment 7h. |
| Ch | Recheck SMART | | Same as segment 4h. |

Figure 192. Short and Extended Self-Test Description

# 10.18 Idle Time Function

The drive periodically saves data in logs and PFA counters in the reserved area of the disks. The information is used by the drive to support various SCSI commands and for the purpose of failure analysis. This is performed if the drive has not received a SCSI command for 10 minutes.

# 10.19 Command Time out Limits

The 'Command Timeout Limits' are defined as the time period from the SCSI Arbitration phase through the SCSI Task complete message, associated with a particular command.

The following times are for environments where Automatic Reallocation is disabled and there are no queued commands.

## 10.19.1 Reassignment Time

The drive should be allowed a minimum of 15 seconds to complete a "Reassign Blocks" command.

## 10.19.2 Format Time

An average of 60 minutes should be allowed to complete a "Format Unit" command. If the vendor unique mode page 00h bit named "FFMT" is set equal to '1'b then the drive should be allowed 30 seconds to complete.

## 10.19.3 Start/Unit Stop Time

The drive should be allowed a minimum of 30 seconds to complete a "Start Stop Unit" command (with Immed bit = 0 ). Initiators should also use this time to allow startup sequences initiated by auto start ups and "Start Stop Unit" commands (with Immed bit = 1) to complete and place the drive in a "ready for use" state.

**Note:** A timeout of one minute or more is recommended but NOT required. The larger system timeout limit allows the system to take advantage of the extensive ERP/DRP that the drive may attempt in order to successfully complete the startup sequence.

## 10.19.4 Medium Access Command Time

The timeout limit for medium access commands that transfer user data or non-user data or both should be a minimum of 30 seconds. These commands are

| | |
|---|---|
| Log Sense | Reserve |
| Mode Select(6) | Rezero Unit |
| Mode Sense(6) | Send Diagnostic |
| Pre-Fech | Write(6) |
| Read(6) | Write(10) |
| Read(10) | Write and Verify |
| Read Capacity | Write Buffer |
| Read Defect Data | Write Long |
| Seek(10) | Write Same |
| Read Long | Verify |
| Release | |

**Note:** The 30-second limit assumes the absence of bus contention and data transfers of 64 blocks or less. This time should be adjusted for anticipated bus contention and if longer user data transfers are requested.

When Automatic Reallocation is enabled, add 45 seconds to the timeout of the following commands;

Read(6)
Read(10)
Write(6)
Write(10)
Write and Verify
Write Same

## 10.19.5 Timeout Limits for Other Commands

The drive should be allowed a minimum of 5 seconds to complete these commands:

Inquiry
Request Sense
Read Buffer
Start/Stop Unit ( with Immed bit = 1 )
Synchronize Cache
Test Unit Ready

The command timeout for a command that is not located at the head of the command queue should be increased by the sum of command timeouts for all of the commands that are performed before it is.

# 10.20 Recommended Initiator ERP

The Drive's design points for error reporting to the system assumes certain system action for the error return codes. These assumptions are:

1. SCSI-2 protocol will be the first priority in reporting errors.
2. The system will maintain a log of all reported errors.
3. System architecture should include all error handling recommendations made in this section. Deviations should have mutual agreement between Drive development and system integration.

This section is directed toward documenting the assumptions made by the Drive that the system is expected to implement. The two error classes that the system should be concerned with are DATA and NON-DATA errors.

Data errors are those errors that deal with the handling of data to and from the MEDIA and are identified by the Additional Sense Code contained in the sense data. The Additional Sense Codes for data errors are as follows:

- OC - Write error
- 11 - Unrecovered read error
- 14 - No record found
- 16 - Data Synchronization mark error
- 17 - Recovered read error without ECC correction
- 18 - Recovered read error with ECC correction

Typically, data errors do not include positioning of the heads or the data path though the electronics.

Nondata errors are those errors that do not have a direct relationship with transferring data to and from the media. Nondata errors can include data handling if the media is not associated with the error (that is, interface error ).

The system action assumed for each class of error is outlined here.

## 10.20.1 Drive Service Strategy

The Drive service strategy is defined so the customer will be able to use the system as soon after a failure is detected as possible. The first priority is to replace the entire drive to make the system operational with minimal service time. The service representative should:

1. Back up all the customer data on this drive if possible
2. Replace the complete drive
3. Restore the customer data
4. Return the drive to customer service

## 10.20.2 Recommendations for System Error Log

The system error log should contain information about the Drive error that will allow recovery actions. The system error logs should contain all the error information returned in the sense data. At a minimum, the following information about each error occurrence should be logged.

- Valid bit and error code (Sense byte 0)
- Sense Key (Sense byte 2)
- Information bytes (Sense bytes 3 through 6)
- Command specific information (Sense bytes 8 through 11)
- Additional Sense Code (Sense byte 12)
- Additional Sense Code Qualifier (Sense byte 13)
- Field Replaceable Unit (Sense byte 14)
- Sense Key Specific (Sense bytes 15, 16, and 17)
- Vender Unique error information (Sense bytes 20 through 23)

## 10.20.3 Data Recovery Procedure

Statistically, most data error activity is noise related and has nothing to do with defects in the media. It is wrong for the system to assume that every data error reported occurred because of a defect in the media. It is also wrong for the system to assume that every data error that occurred because of a media defect rendered the Drive unusable.

Recurring data error activity at the same physical location is an indication of a problem. The problem can be due to a media defect or magnetic damage. A media defect is physical damage to the recording capability of the media while magnetic damage is a defect in the bit pattern written to the media.

In both cases, the error can be corrected without replacing the unit. The physical sector may require relocation. The Drive determines the need to reassign a sector. The Mode Select Page 1 option bit ARRE (See Section "Mode Select") set active allows the Drive to relocate recovered read data errors. Nonrecovered data errors or the ARRE bit being inactive will have additional sense codes returned to recommend reassignment of sectors.

The need to reassign a sector should be infrequent. Sites not meeting error rate criteria are removed from use during SAT (Surface Analysis Test) in Drive manufacturing. With the exception of some early life SAT escapes (sites that were marginally missed during SAT), reassigning defective sectors should be rare. Frequent sector reassignment may be an (early) indication of another type of failure. Sector reassignments are monitored as part of the predictive failure analysis. When a threshold is exceeded, the Drive will notify the initiator that a scheduled service action is required.

Drive soft error rates are based on extraneous random faults that are not predictable. Media defects discovered after the Drive completes manufacturing final test need to be relocated so that soft error rates are not influenced by predictable known error sites. Failure of the system to properly relocate defective media sites can have a direct influence on system throughput and drive error rates.

### 10.20.3 .1 Reassign a Physical Sector

The Drive determines the need to reassign physical sectors based on error activity. Once a physical sector requires reassignment, the Drive will either reassign the physical sector, or recommend to the initiator that the LBA associated with the physical sector be reassigned.

When the following Sense Key, Additional Sense Code, and Additional Sense Code Qualifier combinations are returned, the initiator should reassign the LBA reported at the next opportunity.

**Note:** In Figure 187, the Key, Code, and Qualifier fields are all hex values (i.e., Sense Key 1 is 1h, Sense Code 17 is 17h, etc).

| Key | Code | Qual | Description |
|---|---|---|---|
| 1 | 17 | 07 | Recovered data without ECC - Recommend Reassignment |
| 1 | 18 | 05 | Recovered data with ECC - Recommend Reassignment |

Figure 193. Recommend Reassign Errors

To reassign an LBA that has sense data recommending a reassignment, the initiator should:

1. Attempt to recover the data from the sector being reassigned with a Read (08) or Read (28) command.
2. Reassign the LBA using the Reassign Blocks (07) command.
   - If the reassignment completes successfully (Good Status), log the error in the system error log.
   - If the reassignment completes unsuccessfully (Check Condition Status), follow the "Reassign Blocks Recovery" procedure.
3. Write the LBA that was reassigned.

### 10.20.3 .2 Data Error Logging

The Drive will report data errors to the initiator that do not require immediate action (successful auto reallocation, successful auto rewrite, or no action needed on this occurrence). The initiator should log these errors in the system error log. No other action is required.

| Key | Code | Qual | Description |
|---|---|---|---|
| 1 | 16 | 00 | Data synchronization mark error |
| 1 | 17 | 01 | Recovered data with retries |
| 1 | 17 | 06 | Recovered data without ECC - Auto Reallocated |
| 1 | 17 | 09 | Recovered data without ECC - Data Rewritten |
| 1 | 18 | 00 | Recovered data with ECC |
| 1 | 18 | 02 | Recovered data with ECC - Auto Reallocated |
| 1 | 18 | 07 | Recovered data with ECC - Data Rewritten |

Figure 194. Log Only Errors

### 10.20.3 .3 Reassign Blocks Recovery

The Drive provides the capability to remove media defects without reducing capacity. If the mode parameter bit ARRE is active, the Drive will automatically reallocate LBA's determined to be defective. For those LBA's where the error is unrecoverable or the initiator elects to not have the Drive automatically reallocate LBA's, the Drive will recommend reassignment of the LBA.

Recovery from a failed reassignment consists of the following actions:

- Updating the defect descriptor to remove the LBA's that have been successfully reassigned and then retry the Reassign Blocks command. The LBA contained in the Command Specific Information field of the Sense Data is the LBA in the first defect descriptor that was not reassigned because of the failure. If the command failed because of an unrecoverable read error other than those specified in the defect descriptor, add this LBA to the defect descriptor and retry the command. Refer to Section, "Reassign Blocks" for additional information.
- If the retried Reassign Blocks (07) command completes successfully, returning to normal processing.
- If the retried Reassign Blocks (07) command fails, servicing the drive using the service guidelines recommended in Section 10.20.1 "Drive Service Strategy" on page 221.

## 10.20.4 Nondata Error Recovery Procedure

The Drive will follow a logical recovery procedure for nondata errors. The initiator options for non-data errors are limited to logging the error, retrying the failing command, or replacing the drive.

These recovery procedures assume the initiator practices data back-up and logs errors at the system level for interrogation by service personnel.

### 10.20.4 .1 Drive Busy

The Drive is busy performing an operation. **This is not an error condition.** The initiator can test for completion of the operation by issuing *Test Unit Ready (00)* (or media access) commands.

- If the *Test Unit Ready (00)* (or media access) command completes with *Check Condition Status* then issue a *Request Sense (03)*
  - If the specified recovery procedure for the sense data is for a condition other than drive busy, follow the recovery procedure for the condition reported.
  - If the specified recovery procedure for the sense data is for a drive busy condition, then continue re-issuing the *Test Unit Ready (00)* and *Request Sense* commands for the duration of a media access timeout or until the drive returns *Good Statu*s.

    If the drive has been busy for longer than " Command Time out Limits" (see page 219) specifies, then service the drive using the service guidelines recommended in Section 10.20.1 "Drive Service Strategy" on page 221. Otherwise return to normal processing.
- If the *Test Unit Ready (00)* (or media access) command completes with *Good Statu*s, then return to normal processing.

### 10.20.4 .2 Unrecovered Drive Error

The initiator should retry the failing command.

1. If the retry of the failing command completes with *Good Status* or recovered Sense Key, follow the recovery procedure in Section, "Recovered Drive Error."
2. If the retry of the failing command completes with hardware error sense, verify there is no outside cause (e.g., power supply) for the failure, then retry the failing command.
   i. If the retry of the failing command completes with *Good Statu*s, follow the recovery procedure in next Section, "Recovered Drive Error."
   ii. If the retry of the failing command completes with Recovered sense or Hardware error sense, then service the drive using the service guideline recommended in Section 10.20.1 "Drive Service Strategy" on page 221.

### 10.20.4 .3 Recovered Drive Error

The Initiator should log the error as soft with the recovery level.

### 10.20.4 .4 Drive Not Ready

The initiator should do the following:

1. Issue a *Start Stop Unit (1B)* command.
2. Verify that the drive comes ready within the time specified in "  Command Time out Limits" on page 219.
3. If the drive fails to come ready within the specified time, service the drive using the service guidelines specified in Section 10.20.1  "Drive Service Strategy" on page 221.
4. Retry the failing command.
   i. If the failing command completes with *Good Statu*s, log the error as recovered.
   ii. If the failing command completes with Not Ready sense, verify there is no outside cause (for example, the power supply).  Then service the drive using the service guidelines specified in Section 10.20.1  "Drive Service Strategy" on page 221.

### 10.20.4 .5 No Defect Spare

Three conditions can cause this error:

1. When the *Reassign Blocks (07)* command is issued and there are no spares available for the Drive to use for the relocation requested.
2. When the Glist is full and the sector to be reassigned cannot be added.
3. During a format operation, there was not enough space available to fulfill the spare requirement (Dlist is too large).

Service the Drive following the Section 10.20.1  "Drive Service Strategy" on page 221.

### 10.20.4 .6 Degraded Mode

Refer to 4.1.9 for the definition of this state. There are three causes for entering degraded mode. In all cases the Sense Key is *Not Ready*. The causes are the following:

1. Sense Code/Qualifier of *Logical Unit Not Ready*, *initializing command required*. The spindle motor not spinning or not at the proper speed.
2. This may not be an error condition. The initiator should issue a *Unit start (1B)* command to start the spindle motor. If the Drive fails to come ready in the specified time ( Command Time out Limits " on page 219 ) service the drive using the service guideline recommended in Section 10.20.1  "Drive Service Strategy" on page 221.
3. Sense Code/Qualifier of *Diagnostic Failur*e. Failure of a Send Diagnostic self test, a start up sequence, or other internal target failures.
   - Failure of a send diagnostic self test or a start up sequence.

     This failure is the result of the diagnostics that are executed during power on or when the *Send Diagnostic (1D)* command is executed detecting a failure. As with the RAM code not loaded and the configuration data not loaded, the recovery is either a power cycle or issuing the *Send Diagnostic (1D)* command with the self test bit set active.

Recovery for a failed Send Diagnostic (1D) is achieved in one of the following ways:

- Executing the Send Diagnostic (1D) command
- Power cycling the drive

If the failure repeats, service the drive using the service guideline recommended in Section 10.20.1 "Drive Service Strategy" on page 221.

Recovery for a failed power up sequence is achieved in one of the following ways:

- Issuing a Unit start (1B) command
- Power cycling the drive.

If the failure repeats, service the drive using the service guideline recommended in Section 10.20.1 "Drive Service Strategy" on page 221.

- Internal target failures

  The drive periodically adjusts the track following for each head to compensate for expansion and contraction of the disks due to temperature changes. If one of these adjustments fails, the drive will enter a degraded mode to prevent writing data off track.

  Recovery of this condition is either a power cycle or successful completion of the Send Diagnostic (1D). Service the drive using the recommended service guidelines specified in Section 10.20.1 "Drive Service Strategy" on page 221 if the power cycle or the Send Diagnostic (1D) command fail to complete successfully.

4. Sense Code/Qualifier of **Format Command Failed** Format Unit (04), Sense Code/Qualifier of **Medium Format Corrupted Reassign Failed** Reassign Blocks (07) command, or an automatic reallocation failed or was abnormally terminated.

   Recovery from a failed Format Unit (04) is achieved by retrying the command. If the command fails a second time, service the drive following the procedure defined in Section 10.20.1 "Drive Service Strategy" on page 221.

If the above defined recovery procedures fail to clear the degraded mode condition, the Drive should be replaced. Follow the procedure in Section 10.20.1 "Drive Service Strategy" on page 221 when replacing the drive.

## 10.20.4 .7 Reserved Area Hard Error

Sectors found defective in the reserved area of the disk cannot be reassigned after the Drive leaves the factory. The data in the reserved area is not directly accessible by the initiator. For this reason, the reserved area has all data. A data error must occur in both copies of the data record before the Drive considers a reserved area read error. When this happens, the integrity of the drive is questionable.

Service the Drive using the Section 10.20.1 "Drive Service Strategy" on page 221.

## 10.20.4 .8 Interface Protocol

For all interface protocol errors, the initiator should complete the following steps:

1. Correct the parameter that caused the Illegal Request
2. Retry the failing command
3. If the first retry of the failing command completes with
   - *Good Status*, log the error as recovered

- *Check Condition Status* with sense data for an Illegal Request, verify there is no outside cause (for example, the power supply) for the failure
- *Other*, follow the recommendations for the error condition reported. Retry the failing command. If this retry of the failing command completes with
  - *Good Status*, log the error as recovered
  - *Check Condition Status* with sense data for an Illegal Request, service the drive using the service guideline recommended in Section 10.20.1 "Drive Service Strategy" on page 221.
  - *Other*, follow the recommendations for the error condition reported.

## 10.20.4 .9 Aborted Command

The initiator should determine the cause from the Additional Sense Code (byte 12):

- Sense Key = B (Aborted Command) with Additional Sense Codes of 1B, 25, 43, 49, and 4E are initiator caused abort conditions. The initiator should correct the condition that caused the abort and retry the failing command.
- Sense Key = B (Aborted Command) with Additional Sense Code of 44 or 48 are drive caused abort conditions. The initiator should:
  1. Retry the failing command.
  2. If the retry of the failing command completes with
     - *Good Statu*s, log the error as recovered.
     - Abort Command Sense, verify there is no outside cause (e.g. power supply) for the failure.
  3 Retry the failing command.
  4 If the retry of the failing command completes with
     - *Good Statu*s, log the error as recovered.
     - Abort command sense, then service the drive using the service guideline recommended in Section 10.20.1 "Drive Service Strategy" on page 221.
- Sense Key = B (Aborted Command) and an Additional Sense Code of 47 can be an initiator or Drive caused abort condition. The initiator should follow the above procedure for initiator caused abort conditions if the Drive detected the SCSI bus parity error. The initiator should follow the above procedure for Drive caused abort conditions if the initiator detected the SCSI bus parity error.

## 10.20.4 .10 Unit Attention

Unit Attention Conditions are not errors. They alert the initiator that the drive had an action that may have changed an initiator controlled state in the drive. These conditions are the following:

### Not Ready to Ready Transition

Not ready to ready transition, unit formatted. This *Unit Attention Condition* will not be reported to the initiator that issued the *Format Unit (04*).

### Reset

Reset - This means the drive was reset by either a power-on reset, Bus reset, a Target Reset message, Transceiver Mode Changed reset, or an internal reset.

### Mode Parameters Changed

A *Mode Select (15)* command successfully completed. This means that the mode parameters that are the current value may have changed. The parameters may or may not have changed but the command to

change the parameters successfully completed. The Drive does not actually compare the old current and the new current parameters to determine if the parameters changed. This *Unit Attention Condition* will not be reported to the initiator that issued the *Mode Select (15).*

### Microcode Has Changed

*Write Buffer (3B)* to download microcode has successfully completed. This means that the microcode that controls the Drive has been changed. The code may or may not be the same as the code currently being executed. The Drive does not compare old level code with new code.

### Commands Cleared by Another Initiator

Tagged commands cleared by a clear queue message. This means that the command queue has been cleared. The *Unit Attention Condition* is not reported to the initiator that issued the clear queue message. *Unit Attention Condition* is reported to all initiators that had commands active or queued.

Reissue any outstanding command.

### Log Select Parameters Changed

A Log Select (4C) command successfully completed. This means that the Log Select command cleared statistical information successfully (See Log Select command). Unit Attention Condition is reported to all initiators excluding the initiator that issued the Log Select command.

### Device Identifier Changed

A Set Device Identifier (A4) command successfully completed. This means that the Set Device Identifier information field has been updated. (See Set Device Identifier command.) A Unit Attention Condition is reported to all initiators excluding the initiator that issued the Set Device Identifier command.

## 10.20.4 .11 Components Mismatch

The compatibility test is performed at a power cylcle. The compatibility test verifies the microcode verison of the electronics. When the Drive detects the microcode version mismatch,  the most likely cause is the result of incorrect parts used during a service action.

If the error reported is Key/code/qualifier 4/40/80, Diagnostic failure, bring-up fail, the initiator should do the following:

1. Retry Power cycle
2. Check the send diagnostic end status. If the status is
    - GOOD, Return to normal processing
    - *Check Condition Statu*s, issue a *Request Sense (03)* and follow the recommendations for the sense data returned unless the sense data is for a component mismatch. If the sense data is for component mismatch, service the drive using the service guideline recommended in Section 10.20.1  "Drive Service Strategy" on page 221.

## 10.20.4 .12 Self Initiated Reset

The Drive will initiate a self reset when the condition of the Drive cannot be determined. The internal reset will terminate any outstanding commands, release any reserved initiators, and stop the spindle motor. The initiator can recover by

1. Logging the error
2. Retrying the failing command. If the failing command completes with:
    - *Good Statu*s, return to normal processing
    - Self initiated reset sense, service the drive according the guidelines recommended in Section 10.20.1  "Drive Service Strategy" on page 221.

- Other, follow the recommendations for the error reported.

## 10.20.4 .13 Defect List Recovery

**This is not an error condition**.

The initiator either requested a defect list in a format (block or vendor specific) that the Drive does not support or the requested defect list(s) exceed the maximum list length that can be returned. If the Sense Key/Code/Qualifier are:

1/1F/00, the requested list(s) exceed the maximum length that can be supported. The initiator should request one list at a time. If a single list exceeds the maximum returnable length, this may be an indication of a marginally operational drive. Service the drive following the service guidelines in Section 10.20.1 "Drive Service Strategy" on page 221.

1/1C/01 or 1/1C/02, the requested defect list is not in the format that the Drive supports. The requested defect list is returned in the physical (cylinder, sector, head) format. This is the default format. There is no initiator action required for this condition.

## 10.20.4 .14 Miscompare Recovery

A miscompare can occur on a *Verify (2F)* command or a *Write and Verify (2E)* with the byte check (BytChk) bit active. Recovery for a miscompare error is different for the two commands.

### Verify Command

The initiator should do the following:

1. Verify that the data sent to the drive is the correct data for the byte-by-byte compare.
2. Read the data from the media with a *Read (08)* or *Read (28)* command and verify that the data from the media is the expected data for the byte-by-byte compare.
    - If all data are correct, this is an indication that the data may have been read from the media incorrectly without an error detected. Service the drive using the procedure specified in Section 10.20.1 "Drive Service Strategy" on page 221.
    - If all data are not correct, this is an indication that the data on the media is not the data the initiator expected. Rewrite the correct data to the media.

### Write and Verify Command

The drive uses the same data in the data buffer to write then read and compare. A miscompare error on the *Write and Verify (2E)* command is an indication that the drive cannot reliably write or read the media. Service the drive using the procedures specified in Section 10.20.1 "Drive Service Strategy" on page 221.

## 10.20.4 .15 Microcode Error

The microcode from the interface is validated before the device operates using that microcode. When the validation detects incorrect or incomplete data, the Drive enters degraded mode.

If the initiator attempted to load microcode using the *Write Buffer (3B)* retry the *Write Buffer (3B).* If the command completes with

- *Good Status* - return to normal processing
- *Check Condition Status* - service the drive using the service guidelines recommended in Section 10.20.1 "Drive Service Strategy" on page 221.

If the check sum error occurred during normal processing, the initiator may attempt to load microcode before deciding to service the drive using the service guidelines recommended in Section 10.20.1 "Drive Service Strategy" on page 221.

To load new microcode, the initiator should issue a Write Buffer (3B) command with the download and save option. If the Write Buffer (3B) command completes with

- Good Status, return to normal processing. Retry the failing command. If the task complete with
  - Good Status - Continue normal processing.
  - Check Condition Status for check sum error - Service the drive using the service guidelines recommended in Section 10.20.1 "Drive Service Strategy" on page 221.
  - Check Condition Status for any other error - follow the recommended recovery procedure for the error reported.
- Check Condition Status for Check sum error, service the drive using the service guidelines recommended in Section 10.20.1 "Drive Service Strategy" on page 221.
- Check Condition Status for any other error, follow the recommendations for the returned sense data.

### 10.20.4 .16 Predictive Failure Analysis

The Drive performs error log analysis and will alert the initiator of a potential failure. The initiator should determine if this device is the only device with error activity.

If this drive is the only drive attached to the initiator with error activity, service the drive using the procedures specified in Section 10.20.1 "Drive Service Strategy" on page 221.

**Note**: Service for this drive can be deferred. The longer service is deferred, the more probable a failure can occur that will require immediate service.

If more than this drive is experiencing error activity, the drive is probably not at fault. Locate and service the outside source causing error activity on this drive.

# 11.0 SCSI Sense Data

## 11.1 SCSI Sense Data Format

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

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | **7** | **6** | **5** | **4** | **3** | **2** | **1** | **0** |
| 0 | Valid | Error Code (70h or 71h) | | | | | | |
| 1 | RSVD = 0 | | | | | | | |
| 2 | 0 | | ILI | 0 | Sense Key | | | |
| 3-6 | (MSB) Information Bytes (LSB) | | | | | | | |
| 7 | Additional Sense Length | | | | | | | |
| 8-11 | (MSB) Product Specific Information (LSB) | | | | | | | |
| 12 | Additional Sense Code | | | | | | | |
| 13 | Additional Sense Code Qualifier | | | | | | | |
| 14 | FRU = 0 | | | | | | | |
| 15 | SKSV | Sense-Key Specific Bits | | | | | | |
| 16-17 | Sense-Key Specific Bytes | | | | | | | |
| 18-19 | Reserved = 0 | | | | | | | |
| 20-23 | Vendor unique Error information | | | | | | | |
| 24-29 | Product Specific Information | | | | | | | |
| 30-31 | Reserved = 0 | | | | | | | |

Figure 195. Format of Sense Data.

# 11.2 Sense Data Description

## 11.2.1 Valid (Bit 7 of byte 0)

**0** The Information Bytes (byte 3 through 6) are not defined.

**1** The Information Bytes (byte 3 through 6) contain a valid logical block address.

## 11.2.2 Error Code (Bit 6 - 0 of byte 0)

**70h** Current Error. This indicates an error for the current command.

**71h** Deferred Error. This indicates that the error is for a previous command that has already returned a good status. Such commands are associated with the immediate bit or write caching. Format unit (04h) command is an example of a command that may return a deferred error.

## 11.2.3 ILI: Incorrect Length Indicator (Bit 5 of byte 2)

The ILI bit is valid for the Read Long (3Eh) command and Write Long (3Fh) command only. ILI set to one and Valid Bit set to one indicates that the requested logical block length does not match the logical block length of the data on the medium for a Read Long or Write Long command. The Information field contains residue information about the error. ILL set to zero indicates there is no incorrect length condition.

**0** No Incorrect Length condition.

**1** Incorrect Length Indicated.

## 11.2.4 Sense Key (Bit 3 - 0 of byte 2)

The sense key provides generic categories in which error and exception conditions can be reported. Initiators would typically use sense keys for high level error recovery procedures.

**0h** **No Sense**
There is no sense key information to be reported for the logical unit.

**1h** **Recovered Error**
The last command completed successfully with some recovery action performed by the drive. More detailed information is available in the Additional Sense Code and Additional Sense Code Qualifier.

**2h** **Not Ready**
The logical unit addressed cannot be addressed. More detailed information is available in the Additional Sense Code and Additional Sense Code Qualifier.

**3h** **Medium Error**
The command terminated with an unrecoverable error condition caused by a flaw in the media or an error in the recorded data. More detailed information is contained in the Additional Sense Code and Additional Sense Code Qualifier.

**4h** **Hardware Error**
The drive detected a unrecoverable hardware error while performing a command or during a diagnostic test. More detailed information is contained in the Additional Sense Code and Additional Sense Code Qualifier.

**5h** **Illegal Request**
There was an illegal parameter in the command descriptor block or additional parameter supplied as data. If an invalid parameter is found in the CDB, then the command is terminated without altering the medium. If an invalid parameter is found in parameters supplied as data, then the drive might have altered the medium.

**6h** **Unit Attention**
Indicates that the drive entered in the 'Unit Attention Condition'. (See 11.1.5, “Unit Attention Condition” on page 205.)

**7h-8h** **Not used**

**9h** **Vendor Specific**

**Ah** **Not used**

**Bh** **Aborted command**
The drive aborted the command.

**Ch-Dh** **Not Implemented**

**Eh** **Not Used**

**Fh** **Reserved**

## 11.2.5 Information Bytes (Byte 3 through 6)

This field is only valid when Valid Bit is one.

- ILI = 0 : This field contains the unsigned LBA associated with the sense key. The LBA reported will be within the LBA range of the command as defined in the CDB.

  **Note:** An LBA other than the command LBA may be reported on the Reassign Block (07h) command.

- ILI = 1 : This field contains the difference (residue) of the requested length in bytes. Negative values are indicated by two's complement notation.

## 11.2.6 Additional Sense Length (Byte 7)

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

## 11.2.7 Command Specific Information (Byte 8 through 11)

The values in this field vary with products. Please see the individual product specification for more details.

## 11.2.8 Additional Sense Code/Qualifier (Byte 12 and 13)

The following table shows the description of the combination of Sense Key / Sense Code / Qualifier.

| Key | Code | Qual | Description |
|---|---|---|---|
| 0 | 0 | 0 | No error |
| 0h | 5Dh | 00h | No sense. Predictive Failure Analysis threshold reached |
| 1h | 01h | 00h | Recovered write error no index |
| 1h | 02h | 00h | Recovered no seek comp |
| 1h | 03h | 00h | Recovered write error. Write fault |
| 1h | 0Bh | 01h | Temperature Warning |
| 1h | 0Ch | 01h | Recovered write error with Auto-reallocation. Auto Reallocated |
| 1h | 0Ch | 03h | Recovered write error. Recommend Reassign |
| 1h | 16h | 00h | Recovered write error DAM not found |
| 1h | 17h | 01h | Recovered read error with retries |
| 1h | 17h | 06h | Recovered read error without ECC applied. Auto reallocated |
| 1h | 17h | 07h | Recovered read error without ECC applied. Recommended reassign. This value can be returned only when ARRE = 0 |
| 1h | 17h | 09h | Recovered read error without ECC applied. Data rewritten. This value can be returned only when ARRE = 1 |
| 1h | 18h | 00h | Recovered read error with ECC applied. This value can be returned only when ARRE = 0 |
| 1h | 18h | 02h | Recovered read error with ECC applied. Auto reallocated |
| 1h | 18h | 05h | Recovered read error with ECC applied. Recommended reassign |
| 1h | 18h | 07h | Recovered read error with ECC applied. Data Rewritten |
| 1h | 1Ch | 00h | Defect List Not Found. |
| 1h | 1Ch | 01h | Primary Defect List Not Found. Requested Defect List Format is not supported. Default List Format is returned (Read Defect Data Only) |
| 1h | 1Ch | 02h | Grown Defect List Not Found. Requested Defect List Format is not supported. Default List Format is returned (Read Defect Data Only) |
| 1h | 1Fh | 00h | Partial Defect List Transferred. Defect list longer than 64KB, 64KB of data returned. (Read Defect Data Only) |
| 1h | 44h | 00h | Internal target failure |
| 1h | 81h | 00h | Internal logic error |
| 1h | 5Dh | 00h | Predictive Failure Analysis threshold reached |
| 1h | 5Dh | FFh | Predictive Failure Analysis Test Warning |

Figure 196. Sense Key / Sense Code / Qualifier combinations (part 1 of 4)

| Key | Code | Qual | Description |
|---|---|---|---|
| 2h | 04h | 00h | Not ready. Start spindle motor fail. |
| 2h | 04h | 01h | Not ready. In process of becoming ready. |
| 2h | 04h | 02h | Not ready. Initializing command required. (Unit start) |
| 2h | 04h | 04h | Not ready. Format in progress. |
| 2h | 04h | 09h | Not ready. Self-test in progress. |
| 2h | 31h | 00h | Not ready. Medium format corrupted. |
| 2h | 40h | 80h | Diag Fail - Bring-Up Fail. |
| 2h | 40h | 81h | Diag Fail - HDC Error (**H**ard **D**rive **C**ontroller) |
| 2h | 40h | 85h | Diag Fail - RAM Microcode not Loaded |
| 2h | 40h | 90h | Diag Fail - RRO Calibration Failure |
| 2h | 40h | 91h | Diag Fail - Channel Calibration Failure |
| 2h | 40h | 92h | Diag Fail - Head Load Failure |
| 2h | 40h | 93h | Diag Fail - Write AE Failure |
| 2h | 40h | 94h | Diag Fail - 12V Over Current Failure |
| 2h | 40h | 95h | Diag Fail - Other Spindle Failure. |
| 3h | 03h | 00h | Medium error. Write fault. |
| 3h | 10h | 00h | Medium error. ID CRC error. |
| 3h | 11h | 00h | Medium error. Unrecovered read error. |
| 3h | 16h | 00h | Medium error. Data synchronization mark error. (DAM error) |
| 3h | 19h | 00h | Medium error. Defect list error. A defect list error occurs when a data error is detected while reading the manufacturing defect list or while reading or writing the grown defect list. |
| 3h | 31h | 00h | Medium error. Medium format corrupted. |
| 4h | 01h | 00h | Hardware error. No index or sector. |
| 4h | 02h | 00h | Hardware error. No seek complete. |
| 4h | 03h | 00h | Hardware error. Write fault. |
| 4h | 11h | 00h | Hardware error. Unrecovered read error in reserved area. |
| 4h | 19h | 00h | Hardware error. Defect list error. |
| 4h | 32h | 00h | Hardware error. No defect spare location available. A "no defect spare location available sense code" indicates that the Reassign Block command can not proceed the process because all spare sectors have been used, or it will exceed implementation limitation of defect handling of the drive. |
| 4h | 3Eh | 03h | Hardware error. Self-test failed. |
| 4h | 3Eh | 04h | Hardware error. Unable to update self-test. |

Figure 197. Sense Key / Sense Code / Qualifier combinations (part 2 of 4)

| Key | Code | Qual | Description |
|---|---|---|---|
| 4h | 40h | 80h | Degrade Mode. Diagnostic Fail. Configuration sector valid check fail. Reserved area sector valid check fail. |
| 4h | 40h | 81h | Degrade Mode. Hardware Error |
| 4h | 40h | 85h | Degrade Mode. RAM Microcode Not Loaded. |
| 4h | 44h | 00h | Hardware error. Buffer CRC Error on Write. |
| 4h | 81h | 00h | Hardware error. Internal logic error. |
| 4h | 82h | 00h | Hardware error. Command Time Out. |
| 5h | 1Ah | 00h | Illegal request. Parameter list length error. The number of parameters supplied is not equal to the value expected. |
| 5h | 20h | 00h | Illegal request. Illegal command operation code. This command is also returned when an unsupported command code is received. |
| 5h | 21h | 00h | Illegal request. Logical block address out of range. |
| 5h | 24h | 00h | Illegal request. Invalid field in CDB |
| 5h | 25h | 00h | Illegal request. Invalid lun |
| 5h | 26h | 00h | Illegal request. Invalid fields in the parameter list |
| 5h | 26h | 02h | Parameter value invalid |
| 5h | 26h | 04h | Invalid Release of Persistent Reservation |
| 5h | 2Ch | 00h | Illegal request. Echo buffer being read before being written |
| 5h | 55h | 04h | Illegal request. Insufficient Registration Resources |
| 6h | 28h | 00h | Unit attention. Not ready to ready transition. (Format completed) |
| 6h | 29h | 01h | Unit attention. Power On Reset occurred |
| 6h | 29h | 02h | Unit attention. SCSI Bus Reset occurred |
| 6h | 29h | 03h | Unit attention. Target Reset occurred |
| 6h | 29h | 04h | Unit attention. Self Initiated Reset occurred |
| 6h | 29h | 05h | Unit attention. Transceiver mode change to SE |
| 6h | 29h | 06h | Unit attention. Transceiver mode change to LVD |
| 6h | 2Ah | 01h | Unit attention. Mode Parameters changed |
| 6h | 2Ah | 02h | Unit attention. Log select parameters changed |
| 6h | 2Ah | 03h | Unit attention. Reservations Preempted |
| 6h | 2Ah | 04h | Unit attention. Reservations Released |
| 6h | 2Ah | 05h | Unit attention. Registrations Preempted |
| 6h | 2Fh | 00h | Unit attention. Command cleared by another initiator |
| 6h | 3Fh | 01h | Unit attention. Micro code has been changed |

Figure 198. Sense Key / Sense Code / Qualifier combinations (part 3 of 4)

| Key | Code | Qual | Description |
|---|---|---|---|
| 6h | 3Fh | 05h | Unit attention. Device Identifier Changed |
| 6h | 5Dh | 00h | Predictive Failure Analysis threshold reached |
| 6h | 5Dh | FFh | Predictive Failure Analysis Test Warning |
| Bh | 1Bh | 00h | Aborted command. Synchronous data transfer error. (Extra ack detected) |
| Bh | 25h | 00h | Aborted command. Unsupported LUN. The drive supports LUN 0 only |
| Bh | 3Fh | 0Fh | Aborted command. Echo buffer overwritten |
| Bh | 43h | 00h | Aborted command. Message reject error. A message reject error occurs when an inappropriate or unexpected message reject is received from the initiator or the initiator rejects a message twice |
| Bh | 44h | 00h | Aborted command. Buffer CRC Error on Read |
| Bh | 45h | 00h | Aborted command. Selection/Reselection failed. A selection/reselection error occurs when the initiator fails to respond to a reselection within 250 milli-seconds after the drive starts reselection. The reselection is attempted a second time before setting selection/reselection failed sense code |
| Bh | 47h | 00h | Aborted command. SCSI parity error |
| Bh | 48h | 00h | Aborted command. Initiator detected error message received. An initiator detected error occurs when the initiator detects an error, sends a message to retry, detects the error again, and sends the retry message a second time. The drive then sets check condition status with Initiator Detected Error |
| Bh | 49h | 00h | Aborted command. Inappropriate/illegal message. An inappropriate or illegal message occurs when the initiator sent a message that either is not supported or is not in a logical sequence |
| Bh | 4Eh | 00h | Aborted command. Overlapped commands attempted |

Figure 199. Sense Key / Sense Code / Qualifier combinations (part 4 of 4)

## 11.2.9 FRU : Field Replaceable Unit (Byte 14)

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

## 11.2.10 Sense Key Specific (Byte 15 through 17)

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

### 11.2.10.1 Sense Key Specific - Illegal Request (Sense Key = 5h)

Error field pointer is returned.

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | **7** | **6** | **5** | **4** | **3** | **2** | **1** | **0** |
| 15 | SKSV | C/D | Reserved | | BPV | Bit Pointer | | |
| 16<br>17 | (MSB) Field Pointer (LSB) | | | | | | | |

Figure 200. Field Pointer Bytes

**SKSV** Sense-key specific valid

0 Sense-key specific field is not valid.

1 Sense-key specific field is valid.

**C/D** Command/Data

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

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

**BPV** Bit Pointer Valid

0 Bit pointer field is not valid.

1 Bit pointer field is significant.

**Bit Pointer** Indicates which bit of the byte number reported in Field Pointer is the bit in error. When a multiple bit field is in error, the pointer points to the most significant bit of the field.

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

### 11.2.10.2 Sense Key Specific - Recovered (Sense Key = 1h)

Hardware (Sense Key = 4h) or Medium Error (Sense Key = 3h)

Actual Retry Count is reported.

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
| 15 | SKSV | Reserved | | | | | | |
| 16<br>17 | (MSB) Actual Retry Count (LSB) | | | | | | | |

Figure 201. Actual Retry Count

**SKSV** Sense-key specific valid

0 Actual Retry Count is not valid.

1 Actual Retry Count is valid.

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

### 11.2.10.3 Sense Key Specific - Not Ready (Sense key = 2h)

These fields are defined for the Format unit (04h) command with the Immediate bit set to one and the Send Diagnostic (1Dh) command with Background self-test function.
Progress indication is returned.

| Byte | Bit | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
| 15 | SKSV | Reserved | | | | | | |
| 16<br>17 | (MSB) Progress Indication (LSB) | | | | | | | |

Figure 202. Progress Indication

**SKSV** Sense-key specific valid

0 Progress Indication is not valid.

1 Progress Indication is valid.

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

## 11.2.11 Reserved (Byte 18 through 19)

Reserved fields are filled with zero.

## 11.2.12 Vendor unique error information (Byte 20 through 23)

This field gives detailed information about the error. It contains a unique code which describes where the error was detected and which piece of hardware or microcode detected the error depending on current operation (i.e. Power On Reset/Initialization, Read/Write operation, or SMART Alert).

### 11.2.12.1 Power On Reset/Initialization Error (Sense Key 02h)

While Power On Reset/Initialization is in process, Sense Key 02h "Not Ready" is a possible failure and the Vendor Unique Error information bytes 20 and 21 will contain "POR Step" which indicates where the failure was detected during the microcode processing, and the Vendor Unique Error information bytes 22 and 23 will contain "POR Error" which is the reason for failure.

| Sense Data Byte 20 | Description |
|---|---|
| Bit.7 | Jumper Setting - Disable Auto Spin-Up |
| Bit.6 | Format Corrupted - Degraded Mode |
| Bit.5 | Soft Reset Occurred |
| Bit.4 | RAM ERP Not Ready |
| Bit.3 | Not Used |
| Bit.2 | Not Used |
| Bit.1 | Start Spinning |
| Bit.0 | Initialization by Usage Table |

| Sense Data Byte 21 | Description |
|---|---|
| Bit.7 | RAM Table Validation |
| Bit.6 | Head Test Processing |
| Bit.5 | Loading Read-Only Table |
| Bit.4 | Loading Overlay Code |
| Bit.3 | Loading RAM Code |
| Bit.2 | Starting Spindle Motor |
| Bit.1 | Initiate Microcode Loading |
| Bit.0 | Testing Sector Buffer RAM |

Figure 203. Sense Data Byte 20 and 21 for POR Step

| Sense Data Byte 22 | Description |
|---|---|
| Bit.7 | Not used |
| Bit.6 | Not used |
| Bit.5 | Not used |
| Bit.4 | Not used |
| Bit.3 | Not Used |
| Bit.2 | Not Used |
| Bit.1 | Spindle Motor Spin-Up Failed |
| Bit.0 | EEPROM Data Validation Failed |

Figure 204. Sense Data Byte 22 POR Error

| Sense Date Byte 23 | Description |
|---|---|
| 01h | HDC Test Failure |
| 02h | Sector Buffer RAM Test Failure |
| 03h | EEPROM Read Failure |
| 10h | POR Only Fail |
| 11h | Spin-Up Failure |
| 12h | Reserved Area Table Failure |
| 13h | RAM Code Read Failure |
| 14h | RAM Code Signature Failure |
| 15h | Overlay Code Read Failure |
| 16h | Read-Only Table Read Failure |
| 17h | Head Test Failure |
| 18h | RRO Failure |
| 19h | Channel Calibration Failure |
| 1Ah | Head Load Failure |
| 1Bh | Write AE Failure |
| 1Ch | Over Current 12V Failure |
| 1Dh | Spindle Other Failure |
| 21h | Primary Defect List Read Failure |
| 22h | Grown Defect List Read Failure |
| 23h | Mode Page Parameter Table Read Failure |
| 24h | Persistent Reservation Table Read Failure |
| 25h | SMART Parameter Table Read Failure |
| 26h | Log Parameter Table Read Failure |
| 27h | Usage Table Read Failure |
| 81h | RAM Mapping Failure-Invalid Usage ID |
| 82h | RAM Mapping Failure-Invalid Usage Signature |
| 83h | RAM Mapping Failure-Invalid Table Signature |
| 84h | RAM Mapping Failure-Invalid Table Level |

Figure 205. Sense Data Byte 23 for POR Error

### 11.2.12.2 Read/Write Error (Sense Key 03h/04h)

When an error occurs during a read or write operation, the Vendor Unique Error information will contain the HDC Registers (Status and Error).

| Sense Data Byte 20 | Description (Status) |
|---|---|
| Bit.7 | Sector Count Over |
| Bit.6 | ECC Error on LBA |
| Bit.5 | Drive Error Detected |
| Bit.4 | Uncorrectable Error |
| Bit.3 | Over Symbol |
| Bit.2 | Correctable Error Preparation |
| Bit.1 | Error On |
| Bit.0 | Drive On |

| Sense Data Byte 21 | Description (Status) |
|---|---|
| Bit.7 | Memory on |
| Bit.6 | Drive Complete |
| Bit.5 | ID Find |
| Bit.4 | ID Miss |
| Bit.3 | Drive Busy |
| Bit.2 | End of Track |
| Bit.1 | Expected Sector Flag |
| Bit.0 | Event Handler Busy |

Figure 206. Sense Data Byte 20 and 21 for Read/Write Error

| Sense Data Byte 22 | Description (Status) |
|---|---|
| Bit.7 | N/A |
| Bit.6 | N/A |
| Bit.5 | N/A |
| Bit.4 | N/A |
| Bit.3 | N/A |
| Bit.2 | Sector Pulse Missing |
| Bit.1 | N/A |
| Bit.0 | Write Fault caused by Servo |

| Sense Data Byte 23 | Description (Status) |
|---|---|
| Bit.7 | Overrun |
| Bit.6 | Write Abort |
| Bit.5 | Write Fault |
| Bit.4 | DAM Mismatch |
| Bit.3 | N/A |
| Bit.2 | REQ Error |
| Bit.1 | 11 bytes length of Thermal Asperity |
| Bit.0 | 1 byte length of Thermal Asperity (Read)<br>Channel Fault Detect (Write) |

Figure 207. Sense Data Byte 22 and 23 for Read/Write Error

### 11.2.12.3 SMART Alert (Sense Code 5Dh)

When SMART Alert occurs (Sense Code 5Dh), the Vendor Unique Error information byte 21 will contain Figure 1. the following reason code.

| Sense Data Byte 21 | SMART Alert Reason |
|---|---|
| 00h | Power On Hour Warning |
| 08h | Thermal Sense Warning |
| 14h | Spare Sector Availability Warning |
| 32h | Read Error Rate exceeding the threshold |
| 43h | Seek Error Rate exceeding the threshold |
| 4Ah | Write Error Rate exceeding the threshold |
| 50h | Load/Unload Cycles Count Warning |
| 56h | Spin Up Time Warning |
| 5Bh | Spin Up Retry Count Warning |
| FFh | False Report due to TEST (Mode Page 1C) |

Figure 208. Sense Data Byte 21 for SMART Alert Reason Code

## 11.2.13 Physical Error Record (Byte 24 thru 29)

- ILI = 1 - This field contains zeros.
- ILI = 0 - These bytes contain the physical location of the error in cylinder, head, and sector. Byte 24 and 25 are Cylinder high and cylinder low respectively. Byte 26 is the head number. Byte 28 and 29 are sector high and low respectively.

  Byte 27 is retained for compatibility and will contain the sector number for sector 0 to 254 and will be set to 0FFh for sector numbers greater than 254 or if the sector number is undetermined.

  If the head undetermined, its value is to 0FFh. If the cylinder value is undetermined, byte 24 and 25 are set to 0FFFFh. If the sector number is undetermined, bytes 27, 28, and 29 are all set to 0FFh. If cylinder, head, and sector have no relevance to the error, bytes 24 through 29 will all be set to 0FFh. This field is valid with Sense Key 1, 3 and 4 only.

## 11.2.14 Reserved (Byte 30 through 31)

Reserved fields are filled with zero.

# INDEX

### *N*



### *O*



### *P*



### *R*



### *S*



### *T*



### *U*



### *V*



### *W*